arXiv:1411.6299v2 [cs.CC] 26 Mar 2015

Almost Optimal Pseudorandom Generators for Spherical Caps

Pravesh Kothari* Raghu Meka†

March 30, 2015

Abstract

Halfspaces or linear threshold functions are widely studied in complexity theory, learning theory and algorithm design. In this work we study the natural problem of constructing pseudorandom generators (PRGs) for halfspaces over the sphere, aka spherical caps, which besides being interesting and basic geometric objects, also arise frequently in the analysis of various randomized algorithms (e.g., randomized rounding). We give an explicit PRG which fools spherical caps within error ϵ and has an almost optimal seed-length of $O(\log n + \log(1/\epsilon) \cdot \log\log(1/\epsilon))$. For an inverse-polynomially growing error ϵ, our generator has a seed-length optimal up to a factor of $O(\log\log(n))$. The most efficient PRG previously known (due to Kane [34]) requires a seed-length of $\Omega(\log^{3/2}(n))$ in this setting. We also obtain similar constructions to fool halfspaces with respect to the Gaussian distribution.

Our construction and analysis are significantly different from previous works on PRGs for halfspaces and build on the iterative dimension reduction ideas of [48, 12], the *classical moment problem* from probability theory and explicit constructions of approximate *orthogonal designs* based on the seminal work of Bourgain and Gamburd [9] on expansion in Lie groups.

*University of Texas at Austin, Email: kotpravesh@gmail.com
†Email: raghuvardhan@gmail.com

1 Introduction

A halfspace (a.k.a. a linear threshold function) is a function $f : \mathbb{R}^n \to \{-1, 1\}$ of the form $f(x) = \mathsf{sign}(\sum_{i=1}^n w_i x_i - c) = \mathsf{sign}(\langle w, x\rangle - c)$. Here $w_1, w_2, \ldots, w_n, c$ are reals and $\mathsf{sign}(z) = 1$ for every $z \geq 0$ and -1 otherwise. Halfspaces are a simple class of Boolean functions extensively studied in various contexts in computer science beginning with threshold logic in switching circuits [15, 30, 44, 51, 55], circuits with majority and linear threshold gates in complexity theory [22, 23, 26, 1, 2, 19] and voting and social choice theory [56, 32, 18, 58]. In the recent years, they have been studied extensively in learning theory where learning halfspaces (and functions of a few halfspaces) is arguably the central problem and lies at the core of several important machine learning tools such as the perceptron [53], support vector machines [59] and boosting [20].

In this paper, we study the problem of constructing explicit pseudorandom generators (PRGs) for the class of halfspaces. Constructing PRGs for halfspaces (and more generally, polynomial threshold functions (PTFs)) has been intensively studied in the recent years [8, 16, 17, 35, 36, 34, 37, 49]. In addition to being a natural problem in derandomization, efficient PRGs for halfspaces have concrete applications such as derandomization of the Goemans Williamson algorithm for max cut and deterministic estimation of accuracy of halfspace classifiers in machine learning. Before proceeding, we define PRGs for halfspaces formally[1]:

Definition 1.1 (PRG for Halfspaces w.r.t. a distribution ν)**.** *A PRG for halfspaces with respect to a distribution ν on $\mathbb{R}^n$, with error ϵ, is a function $G : \{0,1\}^r \to \mathbb{R}^n$ such that for every halfspace $h : \mathbb{R}^n \to \{-1, 1\}$,*

$$\left|\mathbb{E}_{x\sim\nu}[h(x)] - \mathbb{E}_{y\sim\{0,1\}^r}[h(G(y))]\right| \leq \epsilon.$$

The parameter r is called the seed-length *of the PRG. G is said to be explicit, if $G(y)$ can be computed in time polynomial in n.*

Alternatively, G is said to "fool halfspaces with error ϵ w.r.t. ν" when it satisfies the definition above. We will skip the explicit reference to ν when the distribution is clear from context. In this work, we focus on the case when ν is one of two natural distributions: the uniform distribution on the sphere $\mathbb{S}^{n-1} \subseteq \mathbb{R}^n$ and the n dimensional spherical Gaussian distribution $\mathcal{N}(0, 1)^n$. Halfspaces are also referred to as *spherical caps* when seen as functions on $\mathbb{S}^{n-1}$. A simple probabilistic argument shows that there *exist* PRGs with seed-length $r = 2\log(n) + 2\log(1/\epsilon) + O(1)$ for both these distributions.

The question of constructing explicit PRGs for spherical caps was first studied in computational geometry in the form of *discrepancy minimization* for spherical caps (cf. [13]). This line of inquiry led to the seminal work of Lubotzky, Philips and Sarnak ([45]) who used *Ramanujan expanders* to construct a PRG for spherical caps over $\mathbb{S}^2$ (in our language) with a seed-length of $3\log(1/\epsilon)+O(1)$. More recently, Diakonikolas et al. [16] showed that bounded independence fools halfspaces w.r.t. the uniform distribution on the Boolean hypercube, giving a PRG with seed-length $O(\log(n)/\epsilon^2)$. Subsequently, Karnin, Rabani and Shpilka [39] developed a PRG for spherical caps with a seed-length of $O(\log n + \log^2(1/\epsilon))$. The best current result due to Kane [34] gives a PRG for spherical caps with a seed-length of $O(\log(n) + \log^{3/2}(1/\epsilon))$.

Despite the long line of works, the seed-length for the best PRGs for halfspaces remains off by poly-logarithmic factors in n for low error regimes (i.e. $\epsilon \approx 1/\mathsf{poly}(n)$). In this work, we resolve

[1] Henceforth, for a multi-set S, $x \sim S$ denotes a uniformly random element of S.

this question for spherical caps and give a construction with seed-length optimal up to a factor of $O(\log\log(n))$.

Theorem 1 (Main Theorem)**.** *There exists a PRG for halfspaces with seed-length* $O(\log(n) + \log(1/\epsilon)\cdot\log\log(1/\epsilon))$ *for error* ϵ*, on the uniform distribution on sphere* $\mathbb{S}^{n-1}$*.*

Our result extends to fool halfspaces with respect to Gaussian distributions as well.

As we describe next, our construction departs significantly from the previous work on constructing PRGs and introduces new ingredients which may be useful elsewhere. In particular, our construction uses an iterative dimension reduction approach as in the works of [48, 12] and makes use of explicit constructions of approximate *orthogonal designs* which are related to quantum analogues of classical k-wise independence and expanders [5, 27, 28, 29, 25, 10]. The analysis of the construction is motivated by the classical *truncated moment problem* from probability theory.

1.1 Outline of constructions

We now give a high-level description of our construction of the PRG and its analysis for spherical caps. For simplicity, we focus on the case of $\epsilon = 1/\mathsf{poly}(n)$ aiming to achieve a seed-length of $O(\log(n)\log\log(n))$. Consider a vector $w \in \mathbb{R}^n$ describing a half space $\mathsf{sign}(\langle w, x\rangle - c)$. Without loss of generality, we can assume that the vector w is normalized so that $||w|| = \sqrt{\sum_{i=1}^n w_i^2} = 1$. It follows immediately from the definitions that to construct a PRG for spherical caps, it suffices to construct a generator $G : \{0,1\}^r \to \mathbb{R}^n$ such that

$$\mathsf{dcdf}(\langle w, G(y)\rangle, \langle w, x\rangle) < \epsilon,$$

where $y \sim \{0,1\}^r$ and $x \sim \mathbb{S}^{n-1}$ and dcdf denotes the CDF or Kolmogorov distance between real-valued random variables.

Let X be the random variable $\langle w, x\rangle$ for $x \sim \mathbb{S}^{n-1}$. We can think of X as obtained by projecting w on to a uniformly random direction. or equivalently, a random one-dimensional *subspace* of $\mathbb{R}^n$. Our construction will exploit this geometric viewpoint by using the following trivial observation:

Fact 1.1. *For any* $1 \le m \le n$*, first picking a random* m*-dimensional subspace* $V \subseteq \mathbb{R}^n$ *and then picking a random one-dimensional subspace of* V *gives the same distribution as picking a random one-dimensional subspace in* $\mathbb{R}^n$*.*

We use the above observation by iteratively projecting the vector w into $\sqrt{n}$ dimensions, and then to $n^{1/4}$ dimensions and so forth until we work in a space of dimension $\Theta(\log n)$. Once we are down to vectors of dimension $\Theta(\log n)$, we use a direct approach to project down to a one-dimensional subspace. We will ensure that each one of these projections can be carried out with $O(\log(n/\epsilon))$ random bits and preserves the properties (including closeness in CDF distance) that we want. Thus, the total randomness used by our generator will be $O(\log(n/\epsilon)\cdot\log\log(n))$ random bits.

To make the above outline concrete let us introduce a central definition[2]:

[2]The use of $\sqrt{n}$ below is somewhat arbitrary and any n^c for $c < 1$ would suffice for us. We choose $\sqrt{n}$ to reduce the number of parameters.

Definition 1.2 (Pseudorandom projections (PRP))**.** *Let Q be a uniformly random projection*[3] *from $\mathbb{R}^n \to \mathbb{R}^{\sqrt{n}}$ and $x \sim \mathbb{S}^{\sqrt{n}}$. A distribution $\mathcal{D}$ on projections from $\mathbb{R}^n \to \mathbb{R}^{\sqrt{n}}$ is an ϵ-pseudorandom projection (PRP) if the following holds for all $w \in \mathbb{R}^n$ and $P \sim \mathcal{D}$:*

$$\mathsf{dcdf}(\langle Pw, x\rangle, \langle Qw, x\rangle) \leq \epsilon.$$

The number of random bits required to sample a P distributed as D is the seed-length *of the PRP.*

Roughly speaking, the above definition says that projecting any vector w to $\sqrt{n}$ dimensions using our PRPs and then projecting to a truly random one-dimensional subspace is indistinguishable from using truly random projections.

Before describing our construction of PRPs let us note how they can be used for constructing PRGs fo spherical caps. As described above, we use our PRPs $O(\log \log n)$ times to project our vector down to $\Theta(\log^2(n))$ dimensions. At this point, we invoke the PRG of Impagliazzo et al. [31] for space bounded machines (that has a seed-length of $O(\log(d) \cdot \log(1/\epsilon))$ for fooling halfspaces in d dimensions with error ϵ). To bound the error we just use a union bound to bound the errors of all projection steps and use Fact 1.1.

In order to extend the construction above for halfspaces on the Gaussian distribution, we observe that $\langle w, g\rangle$ for $g \sim \mathcal{N}(0,1)^n$ is identically distributed as $||g||_2 \cdot \langle w, x\rangle$ for $x \sim \mathbb{S}^{n-1}$ and $||g||_2$, the length of the random (and independent of x) Gaussian vector g (equivalently, a χ-distributed random variable with n degrees of freedom). If G is the PRG described above for spherical caps, we obtain that $\chi_{n,\epsilon} \cdot G$ is a PRG for halfspaces on the Gaussian distribution, where $\chi_{n,\epsilon}$ is obtained by discretization of a χ-distributed random variable with n degrees of freedom.

We next describe our construction of explicit PRPs.

1.2 Pseudorandom projections and the classical moment problem

Note that in our definition of the PRPs, the vector x is truly random over $\mathbb{S}^{\sqrt{n}}$; our analysis will exploit this. At a high level, this helps us because even if our PRPs do not approximate truly random projections very well, the *extra* randomness from x will be enough to mask these imperfections sufficiently.

To describe the intuition more concretely, let us setup some notation. Fix a test vector $w \in \mathbb{R}^n$ and $\mathcal{D}$ be our candidate PRP as in Definition 1.2 and $P \sim \mathcal{D}$. Let Q be a truly random projection from $\mathbb{R}^n$ to $\mathbb{R}^{\sqrt{n}}$. Let $X = \langle Qw, x\rangle$ and $Y = \langle Pw, x\rangle$; our goal is to design $\mathcal{D}$ so that X, Y are close in CDF distance. Our plan to bound $\mathsf{dcdf}(X, Y)$ is to match the low order moments of X and Y. This idea relates to the *classical moment problem*[3]: when do the moments of a (univariate) random variable over a specified range uniquely identify the random variable? In our context, the more relevant question is the *truncated moment problem*: given two real-valued random variables with matching first k moments, how close (under various metrics) are the two random variables?

There is a rich history behind these two questions (see for example, [3], [42]). Unfortunately, the results from the probability literature are quantitatively too weak for us: in most of these general results one needs to match $(1/\epsilon)^{\Omega(1)}$ moments (see [40] and the discussion after Lemma 4.1) to get error ϵ which we cannot afford as we aim for ϵ which is polynomially small.

Our main idea is to exploit the additional structure of the random variables X, Y. For example, the random variable Y above has nicely behaved moments (i.e., not growing too fast) and also has

[3] See Definition 2.2 for a more precise formulation

a smooth, well-behaved (read: bounded derivatives) probability density function (PDF). To exploit this, let us be more concrete. Note that the distribution of random variables X, Y only depend on the norms $\|Pw\|_2, \|Qw\|_2$ respectively (because of the rotational symmetry of $x \sim \mathbb{S}^{\sqrt{n}}$). Thus, if we let $X' = \|Qw\|_2$ and $Y' = \|Pw\|_2$, we can write $X = X' \cdot Z$ and $Y = Y' \cdot Z$, where Z is the random variable obtained by projecting $x \sim \mathbb{S}^{n-1}$ to a fixed direction (say e_1).

It is not hard to see that Z has a smooth pdf and that X' has well-behaved moments (which can be controlled by hypercontractivity). We show that whenever the random variables X', Z satisfy these reasonable conditions, if, in addition, the first k (even order) moments of Y' match the corresponding first k moments of X', then, X, Y are close within an error that is exponentially small in k (the base of the exponent depending on the moments of X' and smoothness of Z). This result fits into the general principle where matching moments with some additional structure can be used to get much stronger quantitative guarantees on closeness of distributions; for example, [14, 6] show similar stronger quantitative bounds for various mixture models.

The above arguments reduce the problem of designing PRPs to that of constructing a distribution $\mathcal{D}$ over projections from $\mathbb{R}^n$ to $\mathbb{R}^{\sqrt{n}}$ such that for $P \sim \mathcal{D}$, and any $w \in \mathbb{R}^n$, the moments of $\|Pw\|_2$ are *approximately* what they should be for a corresponding truly random projection. As it turns out, such distributions have been studied in quantum computing [5, 27, 25, 29, 28, 10, 4] under the label *orthogonal designs*[4]. We discuss them next.

1.3 Orthogonal designs

Orthogonal designs can be seen as generalizations of standard tools in pseudorandomness like k-wise independence and almost k-wise independence to the "uniform" (Haar) distribution over rotation matrices. Let $\mathsf{SO}(n)$ denote all orthogonal matrices in $\mathbb{R}^{n \times n}$ and let $\mathbb{H}$ denote the Haar measure on $\mathsf{SO}(n)$. By polynomials on $\mathsf{SO}(n)$, we mean functions that are polynomials in the entries of matrices from $\mathsf{SO}(n)$.

Definition 1.3 (approximate orthogonal t-design)**.** *A distribution $\mathcal{D}$ on a finite subset of matrices from $SO(n)$ is said to be an ϵ-approximate t-design (in n dimensions) if for every polynomial $p : \mathsf{SO}(n) \to \mathbb{R}$ of degree at most t such that $||p||_1 = 1$ (where $||p||_1$ denotes the sum of absolute values of coefficients of p),*

$$|\mathbb{E}_{\mathcal{D}}[p] - \mathbb{E}_{\mathbb{H}}[p]| \leq \frac{\epsilon}{n^t}.$$

We say that $\mathcal{D}$ is an explicit *orthogonal design if there is a $\mathsf{poly}(n)$ time procedure to sample a matrix according to $\mathcal{D}$. The number of bits of randomness used to sample a matrix according to $\mathcal{D}$ is called its* seed-length*.*

It is not too hard to show using the definitions and the arguments outlined from the previous section, that taking the matrix of first $\sqrt{n}$ rows of an approximate orthogonal t-design one gets a PRP with the same seed-length and error which is ϵ. If we think of fixing the error ϵ (the dimension changes for us as we recurse), to get PRPs with an error of ϵ we need a t-design for $t \approx O\left(\log\left(1/\epsilon\right)/\log\left(n\right)\right)$.

In particular, to get PRPs with nearly-optimal seed-length, it suffices to get approximate t-designs with the near-optimal seed-length. The existence of *finite* orthogonal (or unitary) designs

[4]One usually looks at *unitary designs* in quantum computing, but we ignore this distinction for this high-level discussion.

follows from the general results of Seymour and Zaslavsky [54]. Harrow and Low [28] observe that one can modify the argument of Ambainis et. al. [4] to show that there exist t-designs with optimal (up to constants) parameters. Our application, however, requires efficient *explicit* constructions of these objects and we use the work of Brandao et al. [10] who showed that a recent breakthrough result of Bourgain and Gamburd [9] on expansion in Lie groups implies a construction of approximate orthogonal designs for $t \leq \Theta(n)$. This gives us orthogonal designs with optimal seed-length (up to constant factors).[5]

1.4 Other related work

There is a vast body of work in probability on the *generalized moment problems*, beginning with Stieltjes [57] with a first systematic study appearing in the work of Akhiezer [3] under the name of *classical moment problem*. The ideas are extremely useful in applications in a number of different areas (see the recent textbook of Lasserre [42] for a host of applications). For a survey of various approaches to the moment problem, see the text by Landau [41].

The question of distance (in various metrics) between probability distributions that have (approximately) matching low-degree moments is also well studied, see, for example [33], where the principle measure of distance used is the λ-*metric*. It is possible (see for example [38]) to convert these bounds into the more standard CDF (or Levy) distance bounds using known results [21].

The idea of using stepwise projections in order to reduce the amount of randomness required in each step was successfully employed in constructing almost optimal (with respect to randomness) explicit Johnson-Lindenstrauss (JL) embeddings by Kane et. al. [48]. The analysis in [48] is also based on matching the low-order moments of the lengths of the projections in each step. Their argument, though, is different and simpler as a JL family needs to satisfy only a tail bound condition and one can move from matching low-order moments to tail bounds under simple conditions on the random variables. In contrast, the connection between matching low order moments and CDF distance, as explained above, doesn't hold in general and we crucially exploit the additional smoothening effect of mixing with an independent well behaved random variable to obtain the low errors we need.

Very recently, Gopalan, Kane and Meka [24] gave a PRG for halfspaces whose coefficients are in $\{1, 0, -1\}$ w.r.t the Boolean hypercube with a seed-length of $O((\log(n/\epsilon)) \cdot \mathsf{poly}\log(\log(n/\epsilon)))$; this is incomparable to ours and their methods do not seem to apply in our setting. Their work also uses the iterative dimension reduction approach as in [48] but the actual construction and its analysis are very different from ours.

2 Preliminaries

We start with some notations:

1. For any vector $V = (V_1, V_2, \ldots, V_n) \in \mathbb{R}^n$, $\|V\| = \|V\|_2 = \sqrt{\sum_{i=1}^n V_i^2}$ denotes its Euclidean norm. We will use the same notation for the norm (induced by the inner product) of elements of any infinite dimensional Hilbert space.

[5] As some of the parameters important in our setting are not specified in [10] and we work over real matrices as opposed to Hermitian ones in [10], we give an analysis of the construction of t-designs from the expansion results of Bourgain and Gamburd [9] in Section 6.

2. For any matrix of reals M, $M^\dagger$ denotes its transpose, $\|M\|$ its spectral norm (largest singular value) and $\|M\|_2 = \sum_{i,j} M_{i,j}^2$, its Frobenius (or 2) norm.

3. $\mathcal{N}(0,1)^n$ denotes the spherical multivariate Gaussian distribution in n-dimensions or alternatively, the product Gaussian measure on $\mathbb{R}^n$ with PDF at $X \in \mathbb{R}^n$ given by $(\frac{1}{2\pi})^{-n/2} e^{-\frac{1}{2}\|X\|_2^2}$.

4. χ_n (χ random variable with n-degrees of freedom) denotes the positive real-valued random variable distributed as $Y = \|X\|_2$ where $X \sim \mathcal{N}(0,1)^n$.

5. The Gamma function $\Gamma : \mathbb{R} \to \mathbb{R}$ is defined at any $t \in \mathbb{R}$ by the integral $\Gamma(t) = \int_0^\infty x^{t-1} e^{-x} dx$ and equals $(t-1)!$ whenever t is a positive integer.

Definition 2.1 (CDF or Kolmogorov Distance)**.** *Let X and Y be random variables on some domain D with cumulative distribution functions (CDFs) P_1 and P_2 respectively. The CDF distance between X and Y is defined as* $\mathsf{dcdf}(X,Y) = \sup_{z \in D} |P_1(z) - P_2(z)|$.

2.1 Random rotation and projection matrices

$\mathsf{SO}(n)$ denotes the group (under matrix multiplication) of all real orthogonal $n \times n$ matrices. There is a unique probability measure on $\mathsf{SO}(n)$ invariant under matrix multiplication (on the left or right) by matrices in $\mathsf{SO}(n)$ and is called as the *Haar* distribution (see Section B.1.1 for a brief overview).

$\mathbb{S}^{n-1} \subseteq \mathbb{R}$ denotes the sphere of radius 1 in n dimensions. The uniform distribution on $\mathbb{S}^{n-1}$ is the unique probability distribution on $\mathbb{S}^{n-1}$ invariant under the action of matrices from $\mathsf{SO}(n)$. One can think of $\mathbb{S}^{n-1}$ as the set of all one-dimensional subspaces of $\mathbb{R}^n$ corresponding to each direction (unit vector) it contains. More generally, let $G_{n,t}$ (the *Grassmanian*) denote the set of all t dimensional subspaces of $\mathbb{R}^n$. There exists a unique probability (Haar) measure $\mathbb{H}$ on $G_{n,t}$ such that given any subspace $W \in G_{n,t}$, $\mathbb{H}(W) = \mathbb{H}(O \cdot W)$ where $O \in \mathsf{SO}(n)$ is any rotation matrix. By a uniformly random subspace of t dimensions, we mean an element of $G_{n,t}$ drawn according to the distribution with the PDF $\mathbb{H}$.

To project any given vector $w \in \mathbb{R}^n$ on to a random subspace from $G_{n,t}$, we can draw a matrix Q from $\mathsf{SO}(n)$ distributed according to the Haar measure on $\mathsf{SO}(n)$ and then select the sub matrix formed by the first t rows of Q to obtain R. Then, Rw yields the required random projection. It is a well known fact that the distribution so generated is identical to the Haar measure on $G_{n,t}$.

Definition 2.2 (Uniformly Random Projection Matrix)**.** *Let O be drawn from the Haar distribution on $SO(n)$. For any $m \leq n$ the uniformly random projection matrix from $\mathbb{R}^n$ to $\mathbb{R}^m$ is defined by $Q_{m,n} \in \mathbb{R}^{m \times n}$, the matrix obtained by taking the first m rows of O.*

A standard and useful property of uniforly random projections is that one can perform them stepwise: for any $w \in \mathbb{R}^n$ and $\tilde{m} \leq m \leq n$ let $Q_{m,n}$ and $Q_{\tilde{m},m}$ be independent random projections from $R^n \to \mathbb{R}^m$ and $\mathbb{R}^m \to \mathbb{R}^{\tilde{m}}$. Let $Q_{\tilde{m},n}$ be a uniformly random projection from $\mathbb{R}^n \to \mathbb{R}^{\tilde{m}}$. Then, $Q_{n,\tilde{m}} \cdot w \in \mathbb{R}^{\tilde{m}}$ and $Q_{\tilde{m},m} \cdot Q_{m,n} \cdot w \in \mathbb{R}^{\tilde{m}}$ are identically distributed. We refer the reader to the lecture notes by Vershynin [60] for background on random projections and to the text [11] for background on the orthogonal group.

3 PRGs for spherical caps from pseudorandom projections

In this section, we describe our main result giving a PRG for fooling spherical caps over $\mathbb{S}^{n-1}$ with nearly optimal seed-length:

Theorem 2 (PRG for Spherical Caps)**.** *There exists a PRG for spherical caps on $\mathbb{R}^n$ with error at most ϵ and a seed-length of $O(\log(n) + \log\log(1/\epsilon)\log(1/\epsilon))$.*

We will prove the above result assuming we have constructions of appropriate PRPs as defined in Definition 1.2; we show how to construct PRPs in the subsequent sections.

Theorem 3 (See Section 5)**.** *Fix any $\epsilon > 0$. Then, for any $m = \Omega(\log^2(1/\epsilon))$, there exists an ϵ-PRP from $\mathbb{R}^m \to \mathbb{R}^{\sqrt{m}}$ with a seed-length of $O(\log(m/\epsilon))$.*

When working on small dimensions ($m \sim \log(1/\epsilon)$), we will use the PRG for halfspaces based on the construction for small-space machines due to Nisan [52] and Impagliazzo et. al. [31] as observed in [50].

Fact 3.1. *There exists a PRG $G_{INW} : \{0,1\}^s \to \mathbb{S}^{m-1}$ that fools spherical caps in m dimensions within error ϵ using a seed-length of $s = O(\log(m) \cdot \log(1/\epsilon))$.*

We are now ready to prove Theorem 2. As described in the introduction, the basic idea is to use the PRPs to iteratively reduce the dimension of the space and when the dimension is small enough, we can apply Fact 3.1.

We now describe our generator construction. Fix $\epsilon' = \Theta(\epsilon/\log\log(1/\epsilon)) > 0$ and for $i = 0, 1, 2, \ldots, t-1$ let $n_i = n^{1/2^i}$. Set $t = O(\log\log(1/\epsilon'))$ such that $n_t = \max\{n^{\Theta((1/\log(1/\epsilon')))}, \Theta(\log^2(1/\epsilon'))\}$. Let $\mathcal{D}_i$ denote a PRP from $\mathbb{R}^{n_i} \to \mathbb{R}^{n_{i+1}}$ with error ϵ' and let $G_{INW} : \{0,1\}^s \to \mathbb{S}^{n_t-1}$ be the generator given by Fact 3.1 with error ϵ' and $s = O(\log(n_t) \cdot \log(1/\epsilon'))$. Our generator for fooling spherical caps is defined as follows:

- Sample $P_i \sim \mathcal{D}_i$ for each $i < t$ and let $X_t = G_{INW}(y)$ for $y \sim \{0,1\}^s$, all random draws being independent of each other.
- Output
$$X_0 = (P_0^\dagger P_1^\dagger \cdots P_{t-1}^\dagger X_t) \in \mathbb{S}^{n-1}, \tag{1}$$

Note that because P_i is a projection, the rows of P_i are orthonormal and as $\|X_t\| = 1$ it follows that $\|X_0\| = 1$ so that $X_0 \in \mathbb{S}^{n-1}$. The total seed-length of the generator is the number of bits needed for each of the PRPs and G_{INW}:

$$\sum_{i=1}^{t-1} O(\log(n_i/\epsilon')) + O(\log(n_t) \cdot \log(1/\epsilon')) = O(\log(n) + \log\log(1/\epsilon) \cdot \log(1/\epsilon)). \tag{2}$$

We will show that the above generator fools spherical caps over $\mathbb{S}^{n-1}$ with error $O(\log\log(1/\epsilon) \cdot \epsilon') \leq \epsilon$.

Proof of Theorem 2. We will show that the generator defined by Equation 1 fools spherical caps w.r.t. the uniform distribution on $\mathbb{S}^{n-1}$. The proof is a simple inductive application of the definition of PRPs. We first set up some notation. For $0 \leq i \leq t$, let $Y_i \sim \mathbb{S}^{n_i-1}$ and $X_i = (P_i^\dagger P_{i+1}^\dagger \cdots P_{t-1}^\dagger X_t)$.

Let $Q_i \in \mathbb{R}^{n_{i+1} \times n_i}$ be a truly random projection from $\mathbb{R}^{n_i}$ to a uniformly random n_{i+1}-dimensional subspace inside it. Fooling spherical caps is equivalent to showing that $\mathsf{dcdf}(\langle w, X_0\rangle, \langle w, Y_0\rangle)$ is small. We will show this by induction on $0 \le i \le t$. For brevity, for real-valued random variables Z, Z', we write $Z \approx_\delta Z'$ if $\mathsf{dcdf}(Z, Z') \le \delta$.

Claim 3.1. *For $0 \le i \le t$, and all $v \in \mathbb{R}^{n_i}$,*

$$\mathsf{dcdf}(\langle v, X_i\rangle, \langle v, Y_i\rangle) \le (t - i + 1)\epsilon'.$$

Proof. For $i = t$, the statement in the claim follows from the definition of X_t and Fact 3.1. Suppose the claim is true for $i = j + 1$ for some $0 \le j \le t - 1$. Observe that $X_j = P_j^\dagger X_{j+1}$. Fix $v \in \mathbb{R}^{n_j}$. Now,

$$\langle v, X_j\rangle = \langle v, P_j^\dagger X_{j+1}\rangle = \langle P_j v, X_{j+1}\rangle \approx_{(t-j)\epsilon'} \langle P_j v, Y_{j+1}\rangle \approx'_\epsilon \langle Q_j v, Y_{j+1}\rangle = \langle v, Q_j^\dagger Y_{j+1}\rangle = \langle v, Y_j\rangle,$$

where the first $\approx$ follows from the inductive hypothesis (and the fact that P_j is independent of X_{j+1}) and the second $\approx$ follows from the definition of PRP. Therefore, $\mathsf{dcdf}(\langle v, X_j\rangle, \langle v, Y_j\rangle) \le (t-j+1)\epsilon'$. The claim now follows by induction. □

Therefore, for any $w \in \mathbb{R}^n$, $\mathsf{dcdf}(\langle w, X_0\rangle, \langle w, Y_0\rangle) \le (t + 1)\epsilon' \le \epsilon$. Using the bound on the seed-length from Equation 2, the theorem follows. □

3.1 PRGs for halfspaces with respect to the Gaussian distribution

Next, we describe and analyze the PRG for Halfspaces w.r.t. the spherical Gaussian distribution on $\mathbb{R}^n$. The generator will just be an appropriate scalar multiple of the generator for fooling spherical caps from Theorem 2.

Theorem 4 (PRGs for Halfspaces of Gaussians)**.** *Fix $\epsilon > 0$. There exists a PRG for halfspaces w.r.t. the spherical Gaussian distribution with error at most ϵ and seed-length $s = O(\log n + \log\log(1/\epsilon) \cdot \log(1/\epsilon))$.*

To prove the theorem we shall use the following simple fact.

Lemma 3.1. *For every $n \ge 1$, and $\delta > 0$, there exists a random variable $\chi_{n,\delta}$ samplable efficiently with $O(\log n + \log(1/\delta))$ bits that approximates the random variable χ_n:*

$$\mathsf{dcdf}(\chi_{n,\delta}, \chi_n) \le \delta.$$

Lemma 3.2. *Let U, V_1, V_2 be independent random variables. Then, $\mathsf{dcdf}(U \cdot V_1, U \cdot V_2) \le \mathsf{dcdf}(V_1, V_2)$.*

Proof of Theorem 4. The theorem follows from the following black box reduction and Theorem 2.

Fix $\epsilon' = \epsilon/2$ and let $G : \{0,1\}^s \to \mathbb{S}^{n-1}$ be the generator from previous subsection that ϵ'-fools spherical caps w.r.t. the uniform distribution on $\mathbb{S}^{n-1}$. Let $\chi_{n,\epsilon'}$ denote a ϵ'-approximating random variable for χ_n from Lemma 3.1.

Our generator samples $\chi_{n,\epsilon'}$ and $x \sim \{0,1\}^s$ independently and outputs $\chi_{n,\epsilon} \cdot G(x)$. Let $Y \sim \mathbb{S}^{n-1}$. It is a standard fact that $\mathcal{N}(0,1)^n \equiv \chi_n \cdot Y$ (in law). Now, for any $v \in \mathbb{R}^{n_t}$,

$$\mathsf{dcdf}(\langle \chi_{n,\epsilon'} \cdot G(x), v\rangle, \langle \chi_n \cdot Y, v\rangle) \le \mathsf{dcdf}(\langle \chi_{n,\epsilon'} \cdot G(x), v\rangle, \langle \chi_n \cdot G(x), v\rangle) + \mathsf{dcdf}(\langle \chi_n \cdot G(x), v\rangle, \langle \chi_n \cdot Y, v\rangle) \le$$
$$\mathsf{dcdf}(\chi_{n,\epsilon'}, \chi_n) + \mathsf{dcdf}(\langle G(x), v\rangle, \langle Y, v\rangle) \le 2\epsilon' = \epsilon,$$

using Lemma 3.2. The theorem now follows. □

The objective of the following two sections is to prove Theorem 3.

4 From matching moments to CDF distance

In this section, we give quantitative bounds for the truncated moment problem in terms of the CDF distance: given random variables that have approximately equal low order moments, we derive strong upper bounds on the CDF distance between them. Our bounds are stronger (and crucial to obtaining near optimal seed-lengths for our generators) than those obtained from the general results in probability (see for e.g. [33], [41]) but require stronger analytic properties of the random variables. The results from this section will be used to analyze our construction of PRPs in the next section.

Let Z be a real-valued random variable with mean 0 and CDF $F : \mathbb{R} \to [0,1]$. Our goal here is to show that if X, Y are two non-negative random variables with approximately equal first k moments, then the mixtures $X \cdot Z$ and $Y \cdot Z$ are close in CDF distance. We will show that if F is sufficiently *smooth* and Z satisfies reasonable tail bounds, then the cdf distance between $X \cdot Z$ and $Y \cdot Z$ is exponentially small in k. We first formally define (approximate) moment matching random variables.

Definition 4.1 (Moment-Matching Random Variables)**.** *Random variables X and Y are said to be (k, ϵ)-approximate moment matching if for every polynomial p of degree at most k,*

$$|\mathbb{E}[p(X)] - \mathbb{E}[p(Y)]| \leq \epsilon \cdot \|p\|_1, \tag{3}$$

where $\|p\|_1$ is the sum of absolute values of the coefficients of p.

We are now ready to describe the main technical result of this section:

Lemma 4.1. *Let Z be a mean 0, symmetric random variable with an infinitely differentiable CDF $F : \mathbb{R} \to [0,1]$. Let X, Y be two non-negative $(2k, \epsilon)$-approximate moment matching random variables for even k. Let $\mu^2 = \mathbb{E}[X] \geq 1$ and $\mu_i = \mathbb{E}[\left|X - \mu^2\right|^i]$ for every $i > 1$. Then, for all $\delta \leq 1/2$,*

$$\mathsf{dcdf}(\sqrt{X} \cdot Z, \sqrt{Y} \cdot Z) \leq \delta + 2^{\Theta(k)} \Delta_k \cdot \left(\frac{\mu_k}{\mu^{2k}} + \frac{\sqrt{\mu_{2k}}}{\mu^{2k}} + \epsilon\right),$$

where $\Delta_k = \max\left(1, k^{\Theta(k)} \left(F^{-1}(1-\delta)\right)^k \left(\max_{\ell \leq k} \sup_{2\sqrt{2/3} \leq \frac{z}{F^{-1}(1-\delta)} \leq 2\sqrt{2}} |F^{(\ell)}(z)|\right)\right)$.

It is instructive to compare the error bounds above with the following estimate (due to Klebanov and Mkrtchyan) of CDF distance between random variables that have identical low order moments (the statement for approximately equal low order moment is similar but more cumbersome).

Fact 4.1 (Klebanov and Mkrtchyan [40], Theorem 1 and Remark 2)**.** *Let F, f and G, g be the CDFs and PDFs of real-valued random variables X and Y. Suppose f is bounded and that X, Y have identical finite first $2m$ moments given by $\mu_1, \mu_2, \ldots, \mu_{2k} < \infty$ such that $\mu_2 = 1$. Let $\beta_k = \sum_{i=1}^{k} \mu_{2i}^{1/2i}$. Then, for a universal constant C,*

$$\mathsf{dcdf}(X, Y) \leq C \sup_{t \in \mathbb{R}} f(t) \beta_{k-1}^{-1/4}.$$

We can get a sense of how large k must be to obtain an error smaller than a given ϵ using Fact 4.1: consider the example of $X = \|Qw\|^2$ where Q is a uniformly random projection matrix from $\mathbb{R}^n \to \mathbb{R}^{\sqrt{n}}$. In this case, after scaling X to make $\mu_2 = 1$, we have that $\mu_t^{1/t} = t^{\Omega(1)}$ (follows from an argument similar to the proof of Lemma 5.3). Thus, to obtain an error of ϵ, Fact 4.1 would require matching moments of X and Y of order up to $(1/\epsilon)^{\Omega(1)}$. In contrast, in Section 5, in the proof of Theorem 3 using Lemma 4.1, we will show that for $n = \Omega(\log(1/\epsilon))$, approximately matching only a constant number of moments of X and Y will be enough to give us a CDF distance bound of at most ϵ.

We now move on to the proof of Lemma 4.1. We first collect a few simple facts from elementary analysis that will be useful in our proof of Lemma 4.1. We give proofs for these results in Section A of the Appendix. First we note a bound on the magnitude of derivatives of compositions of infinitely differentiable functions:

Fact 4.2. *Let $g, h : \mathbb{R} \to \mathbb{R}$ be infinitely differentiable functions. Then,*

$$|(g \circ h)^{(k)}(x)| \leq (k!)^2 \cdot \left(\max_{\ell \leq k} |g^{(\ell)}(h(x))|\right) \cdot \left(\max_{\ell \leq k} |h^{(\ell)}(x)|\right)^k.$$

We will also need the following bound on the derivatives of $1/\sqrt{1+x}$ when $x > -1/2$:

Fact 4.3. *For $h : (-1, \infty) \to \mathbb{R}$ be defined by $h(x) = 1/\sqrt{1+x}$. Then, $|h^{(k)}(x)| \leq k!$ for $-1/2 < x$.*

Finally, we write the CDF of a product of two random random variables as a convenient expression:

Fact 4.4. *Let $F : \mathbb{R} \to [0,1]$ be the CDF of a random variable Z. For a positive random variable V independent of Z, the CDF G of $V \cdot Z$ at any $t \in \mathbb{R}$ is given by:*

$$G(t) = \mathbb{E}_V[F(t/V)].$$

We are now ready to prove Lemma 4.1.

Proof of Lemma 4.1. Fix $t \in \mathbb{R}$. Using Fact 4.4, the lemma amounts to obtaining an upper bound on

$$\sup_t \left| \mathbb{E}_X[F(t/\sqrt{X})] - \mathbb{E}_Y[F(t/\sqrt{Y})] \right|.$$

Ideally, we'd like to show that for every t, $F(t/\sqrt{A})$ is well approximated by a low-degree polynomial (in A). We can then invoke the approximate moment matching property of the pair X, Y to bound the difference in the expectations of $F(t/\sqrt{X})$ and $F(t/\sqrt{Y})$. This, however turns out to be too strong. Instead, we will show that F is well approximated by low-degree polynomials *whenever* X, Y do not deviate too far from their expectations. Towards this goal, we first set some notation:

1. $\nu^2 \stackrel{\text{def}}{=} \mathbb{E}[Y]$.
2. $\hat{X} \stackrel{\text{def}}{=} \frac{X-\mu^2}{\mu^2}$, $\hat{Y} \stackrel{\text{def}}{=} \frac{Y-\mu^2}{\mu^2}$.
3. $g = g_t : (-1, \infty) \to \mathbb{R}$ by $g(z) \stackrel{\text{def}}{=} F(\frac{t}{\mu(\sqrt{1+z})})$.

Now, $\mathbb{E}[\hat{X}^k] \leq \frac{\mu_k}{\mu^{2k}}$ and thus, the statement of the theorem is trivial if $\frac{\mu_k}{\mu^{2k}} \geq 1$. Thus, suppose otherwise, in which case $\mathbb{E}[\hat{X}^\ell] \leq 1$ for any $\ell \leq 2k$.

Now, $F(\frac{t}{\sqrt{X}}) = F(\frac{t}{\mu\cdot\sqrt{(1+\hat{X})}}) = g(\hat{X})$ and similarly, $F(\frac{t}{\sqrt{Y}}) = F(\frac{t}{\mu\cdot\sqrt{(1+\hat{Y})}}) = g(\hat{Y})$. Next, we describe the approximating polynomial for g at 0, which will be obtained by truncating the Taylor expansion of g. To bound the error of approximation, we will need to bound the derivatives of g. This can be done whenever $x > -1/2$.

For every $x : |x| < 1/2$, we now apply Fact 4.2 to functions $F(x)$ and $\frac{t}{\mu\sqrt{1+x}}$ and use Fact 4.3 to write:

$$\max_{\ell\leq k}\sup_{|x|<1/2}|g^{(\ell)}(x)| \leq (k!)^3\cdot\left(\frac{t}{\mu}\right)^k\cdot\left(\max_{\ell\leq k}\sup_{-\frac{1}{2}\leq x\leq\frac{1}{2}}\left|F^{(\ell)}\left(\frac{t}{\mu\cdot\sqrt{(1+x)}}\right)\right|\right) \stackrel{\text{def}}{=} \zeta(t).$$

Let $P_k : \mathbb{R}\to\mathbb{R}$ be the degree $k-1$ polynomial obtained by truncating the Taylor expansion of g at 0. Thus, each coefficient of P_k is at most $\zeta(t)$ for $z > -1/2$. Thus, using the error term of the Taylor expansion, we have:

$$|g(z) - P_k(z)| < k\zeta(t)\cdot|z|^k. \tag{4}$$

We can now write the CDF distance between $\sqrt{X}\cdot Z$ and $\sqrt{Y}\cdot Z$ as:

$$\left|\mathbb{E}[F(\frac{t}{\sqrt{X}})] - \mathbb{E}[F(\frac{t}{\sqrt{Y}})]\right| = \left|\mathbb{E}[g(\hat{X})] - \mathbb{E}[g(\hat{Y})]\right| \leq$$
$$\left|\mathbb{E}[g(\hat{X})] - \mathbb{E}[P_k(\hat{X})]\right| + \left|\mathbb{E}[P_k(\hat{X})] - \mathbb{E}[P_k(\hat{Y})]\right| + \left|\mathbb{E}[P_k(\hat{Y})] - \mathbb{E}[g(\hat{Y})]\right|. \tag{5}$$

We will bound each term in the right-side individually. We first use the fact that X and Y are $2k$ moment matching to bound the middle term:

$$\left|\mathbb{E}[P_k(\hat{X})] - \mathbb{E}[P_k(\hat{Y})]\right| \leq \epsilon\cdot\|P\|_1 \leq k\zeta(t)\epsilon. \tag{6}$$

Next, we bound the first term of (5). Let $\mathbb{1}(G_X)$ be the 0-1 indicator of the event $|\hat{X}| < 1/2$ and set $\mathbb{1}(\neg\mathbb{G}_X) = 1 - \mathbb{1}(\mathbb{G}_X)$. Then,

$$|\mathbb{E}[g(\hat{X})] - \mathbb{E}[P_k(\hat{X})]| \leq \mathbb{E}\left[\left|g(\hat{X}) - P_k(\hat{X})\right|\right] \tag{7}$$
$$= \mathbb{E}\left[\left|g(\hat{X}) - P_k(\hat{X})\right|\cdot 1(\mathbb{G}_X)\right] + \mathbb{E}\left[\left|g(\hat{X}) - P_k(\hat{X})\right|\cdot 1(\neg\mathbb{G}_X)\right]$$
When $\mathbb{1}(\mathbb{G}_X) = 1$, we can use Equation 4 to bound the first term (recall that k is even):
$$\leq \mathbb{E}\left[k\zeta(t)\cdot\hat{X}^k\right] + \mathbb{E}\left[\left|g(\hat{X}) - P_k(\hat{X})\right|\cdot 1(\neg\mathbb{G}_X)\right] \tag{8}$$
Applying Cauchy-Schwartz to the second term:
$$\leq \mathbb{E}\left[k\zeta(t)\cdot\hat{X}^k\right] + \mathbb{E}\left[\left|g(\hat{X}) - P_k(\hat{X})\right|^2\right]^{1/2}\cdot\mathbb{E}[1(\neg G_X)]^{1/2}$$
$$\leq k\zeta(t)\cdot\frac{\mu_k}{\mu^{2k}} + 2\left(1 + \mathbb{E}[P_k(\hat{X})^2]^{1/2}\right)\cdot\mathbb{E}[1(\neg(G_X))]^{1/2}. \tag{9}$$

Next, we bound the last two terms of (9). Using the bound on coefficients of P_k and the upper bound of 1 on the moments of $\hat{X}$:

$$\mathbb{E}[P_k(\hat{X})^2]^{1/2} \leq \zeta(t) \cdot \left(\sum_{\ell=0}^{k-1} \mathbb{E}[\hat{X}^{2\ell}]^{1/2} \right) \leq k\zeta(t). \tag{10}$$

By Markov's inequality applied to $\hat{X}^{2k}$,

$$\mathbb{E}[\mathbb{1}(\neg \mathbb{G}_X)] = \mathsf{Pr}[|\hat{X}| \geq 1/2] \leq 2^{2k} \cdot \frac{\mu_{2k}}{\mu^{4k}}. \tag{11}$$

Combining the above bounds we get,

$$\mathbb{E}\left[\left| g(\hat{X}) - P_k(\hat{X}) \right| \right] \leq k\zeta(t) \frac{\mu_k}{\mu^{2k}} + (1 + k\zeta(t)) 2^{k+1} \frac{\sqrt{\mu_{2k}}}{\mu^{2k}}.$$

We now argue similarly for the case of Y. Let $\mathbb{1}(\mathbb{G}_Y)$ be 0-1 indicator of the event $|\hat{Y}| < 1/2$ and $\mathbb{1}(\neg \mathbb{G}_Y) = 1 - \mathbb{1}(\mathbb{G}_Y)$. Now, $\mathbb{E}[\hat{Y}^k] \leq \mathbb{E}[\hat{X}^k] + 2\epsilon \leq \frac{\mu_k}{\mu^{2k}} + 2\epsilon$. Next, as above, an application of Markov's inequality combined with the observation that X, Y are moment matching (and thus $\mathbb{E}[(Y - \mu^2)^{2k}] \leq \mathbb{E}[(X - \mu^2)^{2k}] + 2\epsilon\mu^{4k}$) yields

$$\mathbb{E}[\mathbb{1}(\neg G_Y)] \leq 2^{2k} \cdot \frac{\mathbb{E}[(Y - \mu^2)^{2k}]}{\mu^{4k}} \leq 2^{2k} \frac{\mu_{2k}}{\mu^{4k}} + 2^{2k+1}\epsilon.$$

Thus, arguing as in the case of (9), we obtain:

$$\mathbb{E}\left[\left| g(\hat{Y}) - P_k(\hat{Y}) \right| \right] \leq k\zeta(t) \frac{\mu_k}{\mu^{2k}} + 2k\zeta(t)\epsilon + 2^k(1 + k\zeta(t)) \cdot \frac{\sqrt{\mu_{2k}}}{\mu^{2k}} + (1 + k\zeta(t)) 2^{k+2}\epsilon.$$

Thus, from (5), we finally have:

$$\left| \mathbb{E}[F(t/\sqrt{X})] - \mathbb{E}[F(t/\sqrt{Y})] \right| \leq 2k \left(\zeta(t) \frac{\mu_k}{\mu^{2k}} \right) + (1 + k\zeta(t)) 2^{k+1} \frac{\sqrt{\mu_{2k}}}{\mu^{2k}} + \epsilon \cdot 2^{3k} \zeta(t). \tag{12}$$

When $t \geq 2\mu F^{-1}(1 - \delta)$, it is easy to bound the CDF distance:

$$\begin{aligned}
\mathbb{E}[F(t/\sqrt{X})] &\geq \mathbb{E}[F(t/\sqrt{X}) | \mathbb{1}_{\mathbb{G}_X} = 1] \cdot \mathsf{Pr}[1(\mathbb{G}_X) = 1] \\
&\geq F(t/2\mu) \cdot (1 - \mathsf{Pr}[1(\neg \mathbb{G}_X) = 1]) \geq 1 - \delta - \mathsf{Pr}[1(\neg \mathbb{G}_X) = 1] \\
&\geq 1 - \delta - 2^{2k} \frac{\mu_{2k}}{\mu^{4k}}.
\end{aligned}$$

Similarly, in this case:

$$\mathbb{E}[F(t/\sqrt{Y})] \geq 1 - \delta - 2^{2k} \frac{\mu_{2k}}{\mu^{4k}} - 2^{2k+1}\epsilon.$$

Thus, when $t \geq 2\mu F^{-1}(1 - \delta)$,

$$\left| \mathbb{E}[F(t/\sqrt{X})] - \mathbb{E}[F(t/\sqrt{Y})] \right| \leq \delta + 2^{2k+1} \cdot \frac{\sqrt{\mu_{2k}}}{\mu^{4k}} + 2^{2k+1}\epsilon.$$

On the other hand, when $t \leq 2\mu F^{-1}(1 - \delta)$ (and $|x| \leq 1/2$) yields $\frac{t}{\mu\sqrt{1+x}} \in [2\sqrt{\frac{2}{3}}, 2\sqrt{2}]$ and thus: $\zeta(t) \leq \Delta_k$. We can now use (12) setting $\zeta(t) \leq \Delta_k$ to obtain the lemma. □

5 PRPs from approximate orthogonal designs

In this section, we show how to construct PRPs from approximate orthogonal designs and thus proving Theorem 3, which we first restate here.

Theorem 5 (Theorem 3 restated)**.** *Fix any $\epsilon > 0$. Then, for any $m = \Omega(\log^2(1/\epsilon))$, there exists an ϵ-PRP from $\mathbb{R}^m \to \mathbb{R}^{\sqrt{m}}$ with a seed-length of $O(\log(m/\epsilon))$.*

For this, we shall assume the existence of a good explicit approximate orthogonal design, a proof of which is provided in the next section:

Lemma 5.1 ([10]; see Section 6)**.** *There exists an efficiently samplable ϵ-approximate orthogonal t-design with seed-length $O(t\log(n) + \log(1/\epsilon))$.*

We first describe the idea behind the proof of Theorem 3. Let $\mathcal{D}$ be an ϵ approximate orthogonal $2k$-design in m dimensions and denote $\sqrt{m}$ by $\tilde{m}$. For every $P' \sim \mathcal{D}$, we define $P \in \mathbb{R}^{\tilde{m}\times m}$ to be the matrix obtained by taking the first $\tilde{m}$ rows of P'. Let the resulting distribution on matrices in $\mathbb{R}^{\tilde{m}\times m}$ be denoted by $\mathcal{D}_{\tilde{m}}$. We will show that $\mathcal{D}_{\tilde{m}}$ is a β-PRP where β depends on the parameters k and ϵ. Suppose Q is the uniformly random projection from $\mathbb{R}^m \to \mathbb{R}^{\tilde{m}}$, $w \in \mathbb{R}^m$ satisfies $\|w\| = 1$ and $v \sim \mathbb{S}^{\tilde{m}-1}$. Then, to show that P is an PRP, we must show that $\mathsf{dcdf}(\langle Pw, v\rangle, \langle Qw, v\rangle)$ is small.

Let $Y = \|Pw\|^2$. The random variable $\langle Pw, v\rangle$ has the same distribution as $\sqrt{Y}\cdot\langle w', v\rangle$ for $w'{:}\|w'\| = 1$ and $v \sim \mathbb{S}^{\tilde{m}-1}$ independent of Y. Thus, the random variable of our interest is a mixture $\sqrt{Y}\cdot Z$ where Z is distributed as $\langle w', v\rangle$. Similarly, $\langle Qw, v\rangle$ is distributed as a mixture $\sqrt{X}\cdot Z$ where $X = \|Qw\|^2$.

Thus, if we show that a) Z has an infinitely differentiable CDF with a reasonably bounded tail and derivatives b)X has sufficiently slow growing moments and c) X, Y are approximately moment matching, then we can apply the result from the previous section to show that the CDF distance between $\sqrt{X}\cdot Z$ and $\sqrt{Y}\cdot Z$ is small. In the following, we implement this plan and show that the X, Y, Z defined above indeed satisfy the conditions required to complete the proof of Theorem 3. We first record the required properties of Z:

Lemma 5.2. *Let Z be the random variable distributed as $\langle w, v\rangle$ for $v \sim \mathbb{S}^{m-1}$ and any fixed $w{:}\|w\| = 1$ as above. Let f, F be the PDF and CDF of Z, respectively. Then,*

1. ***Sharp Tail:*** *$F^{-1}(1-\delta) < 1/10$ for $\delta = 0.995^{\sqrt{m}}$.*
2. ***Bounds on the Derivatives:*** *For any $0 < x < 1 : |f^{(q)}(x)| \le c\cdot 10^q \cdot q!/|x|^q$, where $c = \frac{1}{\sqrt{\pi}}\cdot\frac{\Gamma(m/2)}{\Gamma(\frac{m-1}{2})}$.*

Next, we bound the moments of X:

Lemma 5.3 (Moments of X)**.** *Fix any $w \in \mathbb{R}^m$ such that $\|w\| = 1$ and define random variables $X = \|Qw\|^2$ and $Y = \|Pw\|^2$ where Q is the uniformly random projection from $\mathbb{R}^m \to \mathbb{R}^{\tilde{m}}$ for $\tilde{m} = \sqrt{m}$ and $P \sim \mathcal{D}$. Then, for any $p \le \tilde{m}/20$,*

$$\mathbb{E}\left[\frac{|X - \mathbb{E}[X]|^p}{\mathbb{E}[X]^p}\right] \le p^{O(p)}\cdot\tilde{m}^{-4p/5}.$$

We can now complete the proof of Theorem 3 using Lemma 4.1, Lemma 5.2 and Lemma 5.3.

Proof of Theorem 3. Let $k < \tilde{m}/80$ be such that $(\tilde{m}/k)^{-\Theta(k)} < \epsilon$. Since $m = \Omega(\log^2(1/\epsilon))$, such a k exists. Let $\mathcal{D}$ be an ϵ-approximate $4k$-orthogonal design in m dimensions and $\mathcal{D}_{\tilde{m}}$ be the distribution over matrices in $\mathbb{R}^{\tilde{m}\times m}$ obtained by taking the submatrix of first $\tilde{m}$ rows of a random draw from $\mathcal{D}$. Let Q be a uniformly random projection matrix obtained by taking the first $\tilde{m}$ rows of a Haar distributed matrix from $\mathsf{SO}(m)$.

We claim that the map: $w \in \mathbb{R}^m \to Pw \in \mathbb{R}^{\tilde{m}}$ where $P \sim \mathcal{D}_{\tilde{m}}$ is an ϵ-PRP. We set $X = \|Qw\|^2$, $Y = \|Pw\|^2$ and Z distributed as $\langle w, v\rangle$ for $v \sim \mathbb{S}^{\tilde{m}-1}$. Observe that by our construction of P and Q above, each of X, Y can be seen as polynomial of degree 2 on $\mathsf{SO}(m)$. Thus, X^t and Y^t are are polynomials on $\mathsf{SO}(m)$ of degree at most $2k$ for any $t \le k$. It is easy to observe (using that $\sum_i w_i^2 = 1$) that the absolute value of the coefficients of the coefficients of these polynomials is at most m^{2k}. Further, since $\mathcal{D}$ is an ϵ-approximate $4k$-design, we must have, for any univariate polynomial p of at most k and sum of absolute value of its coefficients $||p||_1 = 1$:

$$|\mathbb{E}[p(X)] - \mathbb{E}[p(Y)]| \le \frac{\epsilon}{m^{4k}} \cdot m^{2k} \le \epsilon \cdot m^{-2k}.$$

We use Lemma 5.3 and Lemma 5.2 to obtain the estimates of all parameters required to apply Lemma 4.1 to X and Y above now. Let F and f be the CDF and PDF of Z respectively. For X, Y and Z above, we have, $\Delta_k = k^{\Theta(k)}$ and μ_k/μ^{2k}, $\sqrt{\mu_{2k}}/\mu^{2k} \le m^{-\Theta(k)}$. We set $\delta = 0.995^{\sqrt{m}}$ and note that $F^{-1}(1-\delta) < 1/10$. It is easy to verify that with this setting of the parameters, Lemma 4.1 gives an error bound of ϵ for appropriate setting of constants hidden in the Θs.

Next, we verify the seed-length used for the construction above: observe that the k chosen above can be written as $\max\{\Theta(1), \Theta(\log(1/\epsilon)/\log(m))\}$. Thus, the required seed-length is given by: $O(k\log(m) + \log(1/\epsilon)) = O(\log(m/\epsilon))$ as promised. □

In the remaining part of this section, we prove Lemma 5.2 and Lemma 5.3.

Proof of Lemma 5.2. Because of rotation invariance of $v \sim \mathbb{S}^{\tilde{m}-1}$, Z is identically distributed as the first coordinate of v: v_1. The PDF $f : [-1,1] \to [0,1]$ of Z at any t is given by (we give a proof in the Section A.2 of the Appendix):

$$f(t) = c_{\tilde{m}} \cdot (1-t^2)^{\frac{\tilde{m}-3}{2}}$$

where $c_{\tilde{m}} = \frac{1}{\sqrt{\pi}} \cdot \frac{\Gamma(\tilde{m}/2)}{\Gamma(\frac{\tilde{m}-1}{2})}$. From the expression for f, it is easy to verify that for $\delta = 0.995^{\sqrt{m}}$, $F^{-1}(1-\delta) < 1/10$.

To bound the derivatives of f, we will need the following standard theorem from complex analysis (due to Cauchy):

Fact 5.1 (Cauchy's Estimate)**.** *Let $\eta : \mathbb{C} \to \mathbb{C}$ be a holomorphic function on $D(z,r)$, the closed disk of radius r centered at $z \in \mathbb{C}$, and let $M = \max_{y:|z-y|=r} |F(y)|$. Then,*

$$|\eta^{(k)}(z)| \le \frac{k!}{r^k} \cdot M.$$

We can now estimate the derivatives of f. We first extend f to all of the complex plane to obtain $\tilde{f}$. Then, $\tilde{f}(z) = c_{\tilde{m}} \cdot (1-z^2)^{\frac{\tilde{m}-3}{2}}$ is a polynomial in z and thus holomorphic everywhere and in particular, on every closed disk in the complex plane. The complex derivatives of $\tilde{f}$ at any point x on the real line are equal to the derivatives of f at x. Thus, we can apply Cauchy's estimate to produce an upper bound on the derivatives.

Consider the disk $D(x,r)$ of radius $r = x/10$ centered at a (real) x in the complex plane. We need to estimate the maximum absolute value of $\tilde{f}$ over the boundary of $D(x,r)$. Parametrize the boundary of $D(x,r)$ as $z = x + rcos(\theta) + irsin(\theta)$ for $\theta \in [0, 2\pi]$. We can write $|1 - z^2| = 1 - 2Re(z^2) + |z|^4$. Now, $Re(z^2) = (x + rcos(\theta))^2 - r^2sin^2(\theta) \geq 0.8x^2$. On the other hand, $|z|^4 = (x + rcos(\theta))^2 + r^2sin^2(\theta))^2 \leq (x + r)^4 \leq 1.21x^4$. Thus, $-2Re(z^2) + \|z\|^4 \leq -1.6x^2 + 1.21x^4 \leq 0$. Thus, $|1 - z^2| \leq 1$ for every z on the boundary of $D(x,r)$. Hence, the maximum absolute value of $\tilde{f}$ on $D(x,r)$ is at most $c_{\tilde{m}}$. Applying Cauchy's estimate to $\tilde{f}$ now gives:

$$|f^{(q)}(x)| \leq c_{\tilde{m}} \cdot 10^q \cdot q!/|x|^q.$$

□

Before moving on to prove Lemma 5.3, we collect three facts useful in the proof:

We will need the following Marcinkiewicz-Zygmund inequality for moments and the standard gaussian concentration:

Fact 5.2 (Marcinkiewicz-Zygmund [47] (see also [46]))**.** *Let $S_1, S_2, \ldots, S_q$ be a sequence of independent, zero mean random variables. Then:*

$$\mathbb{E}[|\sum_{i=1}^{q} S_i|^p] \leq p^{O(p)} \cdot q^{p/2-1} \sum_{i=1}^{q} \mathbb{E}[|S_i|^p].$$

The following concentration bound is standard for Gaussian random vectors.

Fact 5.3 (Gaussian Concentration [43] (see also Corollary 2.3 on Page 6 of [7]))**.** *Let $g \in \mathbb{R}^m$ be a vector with each coordinate distributed independently as $\mathcal{N}(0,1)$. Then, for any $0 < \xi < 1$:*

$$\Pr[|\|g\|^2 - m| > \xi \cdot m] \leq 2e^{-\xi^2 \cdot m/4}.$$

Proof of Lemma 5.3. Due to the rotation invariance of the uniformly random projection Q, $X = \|Qw\|^2$ has the same distribution as the Euclidean norm of the vector formed by the first $\tilde{m} = \sqrt{m}$ coordinates of $v \sim \mathbb{S}^{m-1}$. Thus, $\mathbb{E}[X] = \tilde{m}/m = 1/\sqrt{m}$. Recall that v has the same distribution as the random variable $g/\|g\|$ where $g \in \mathbb{R}^m$ has each coordinate distributed $\mathcal{N}(0,1)$. Let $\mathbb{1}_{\mathbb{G}}$ be the

indicator of the event that $m(1-\xi) \leq \|g\|^2 \leq (1+\xi)m$ and $\neg \mathbb{1}_{\mathbb{G}}$, its negation. Then, we have:

$$\mathbb{E}[|X - \mathbb{E}[X]|^p] = \mathbb{E}[|\sum_{i=1}^{\tilde{m}} v_i^2 - \tilde{m}/m|^p] = \mathbb{E}[\frac{|\sum_{i=1}^{\tilde{m}} g_i^2 - \|g\|^2 \cdot \frac{\tilde{m}}{m}|^p}{\|g\|^{2p}}]$$
$$= \mathbb{E}[\frac{|\sum_{i=1}^{\tilde{m}} g_i^2 - \|g\|^2 \cdot \frac{\tilde{m}}{m}|^p}{\|g\|^{2p}} \cdot \mathbb{1}_{\mathbb{G}}] + \mathbb{E}[\frac{|\sum_{i=1}^{\tilde{m}} g_i^2 - \|g\|^2 \cdot \frac{\tilde{m}}{m}|^{2p}}{\|g\|^{2p}} \cdot \neg \mathbb{1}_{\mathbb{G}}].$$

Estimating first term using the bound on $\|g\|^2$

and applying Cauchy-Schwartz to the second term,

$$\leq (\frac{1}{1-\xi})^p \cdot \mathbb{E}\left[\frac{2^p\left(\left|\sum_{i=1}^{\tilde{m}} g_i^2 - \tilde{m}\right|^p + \xi^p \cdot \tilde{m}^p\right)}{m^p}\right] + \mathbb{E}[\neg \mathbb{1}_{\mathbb{G}}]^{1/2} \cdot \left(\mathbb{E}\left[\left|\sum_{i=1}^{\tilde{m}} v_i^2 - \tilde{m}/m\right|^{2p}\right]\right)^{1/2}.$$

Using that $\mathbb{E}[v_i^2] = 1/m$ for each i

$$\leq (\frac{2}{m(1-\xi)})^p \cdot (\xi^p \tilde{m}^p + \mathbb{E}[|\sum_{i=1}^{\tilde{m}} (g_i^2 - 1)|^{2p}]) + 2^p e^{-\xi^2 m/8} \cdot (\frac{\tilde{m}}{m})^p.$$

The first term is easy to estimate using Fact 5.2 (the same result can also be obtained from Bernstein-like inequalities for exponential random variables [61]). We have:

$$\mathbb{E}[|\sum_{i=1}^{\tilde{m}} (g_i^2 - 1)|^p] \leq p^{O(p)} \cdot \tilde{m}^{p/2} \cdot \mathbb{E}[|g_i^2 - 1|^p] = p^{O(p)} \cdot \tilde{m}^{p/2}.$$

Choosing $\xi = 1/m^{1/5}$ and noting that $p < \tilde{m}/20$, we can now obtain: $\mathbb{E}[|X - \mathbb{E}[X]|^p]/(\mathbb{E}[X])^p \leq p^{O(p)} \cdot m^{-2p/5} = p^{O(p)} \cdot \tilde{m}^{-4p/5}$.

□

6 Constructing approximte orthogonal designs

In this section, we give a proof of a construction of approximate orthogonal designs, proving Lemma 5.1, based on a recent result of Bourgain and Gamburd [9]. This also follows from the work Brandao et. al. [10] (Page 17, Equation B2) except for some technicalities and concrete quantitative bounds which we need and work out next. We first provide some background before stating the result of [9] (see the text by Bump [11] for a detailed exposition).

Bourgain and Gamburd[9] show that there exist Cayley graph expanders on $\mathsf{SU}(n)$. This also implies that there exist Cayley graph expanders with the same parameters on the group $\mathsf{SO}(n)$ (see Appendix B). In this section, we use the construction for $\mathsf{SO}(n)$ to obtain approximate orthogonal t-designs. We provide the necessary background and the deferred proofs from this section in Appendix B.

We briefly recall Cayley graphs on finite groups before working on $\mathsf{SO}(n)$. A Cayley graph on a group G is defined by a set of generators (inverse closed) $g_1, g_2, \ldots g_k$. The vertex set is given by the elements of the group G and there is an edge between h, h' iff $h = g_i h'$ for some generator g_i.

One can define a Cayley graph on an infinite group similarly and in the following, we adopt the linear operator view.

Let $\mathcal{L}^2(\mathsf{SO}(n)) = \{\rho : \mathsf{SO}(n) \to \mathbb{R} \mid \int_{\mathbb{H}} \rho^2 < \infty\}$ be the linear space of square integrable functions on $\mathsf{SO}(n)$ with respect to the Haar measure $\mathbb{H}$. For any $g \in \mathsf{SO}(n)$, let $\mathscr{T}_g : \mathcal{L}^2(\mathsf{SO}(n)) \to \mathcal{L}^2(\mathsf{SO}(n))$ be the linear operator on $\mathsf{SO}(n)$ defined so that $\rho \to \mathscr{T}_g \cdot \rho$, where $(\mathscr{T}_g \cdot \rho)(y) = \rho(y \cdot g)$.

Next, we define Hecke (or averaging) operators on $\mathcal{L}^2(\mathsf{SO}(n))$ that correspond to the finite dimensional normalized adjacency matrices described above. We say a set $Gen = \{g_1, g_2, \ldots, g_k\}$ is inverse closed if for every $g \in Gen$, $g^{-1} \in G$.

Definition 6.1 (Hecke Operators)**.** *For some universal constant $k > 0$, an averaging operator (also known as* Hecke *operator) with an inverse closed set of generators $g_1, g_2, \ldots, g_k \in G$ is a linear operator $\mathscr{T} : \mathcal{L}^2(G) \to \mathcal{L}^2(G)$ defined by $\mathscr{T} \stackrel{def}{=} \frac{1}{k} \sum_{i=1}^{k} \mathscr{T}_{g_i}$.*

It is easy to verify that a Hecke operator $\mathscr{T}$ on $\mathcal{L}^2(\mathsf{SO}(n))$ is bounded and compact and thus has a spectrum. Thus we can look at the gap between the first and second eigenvalues of $\mathscr{T}$ to talk of the expansion of the associated graph. This is encapsulated in the following definition:

Definition 6.2 (Hecke Operators with Spectral Gap)**.** *A Hecke operator on $\mathcal{L}^2(G)$ for a compact Lie group G is said to have a spectral gap, if for every $\rho \in \mathcal{L}^2(G)$ such that $\|\rho\|_2 = 1$ and $\mathbb{E}_{g \sim \mathbb{H}}[\rho(g)] = \int_{\mathbb{H}} \rho dg = 0$, $\|\mathscr{T}\rho\|_2 \leq \lambda < 1 - \delta$ for a universal constant $\delta > 0$.*

We can now describe the (consequence of) result of Bourgain-Gamburd [9] that we need in the language of Hecke operators:

Corollary 6.1 (See Corollary B.1)**.** *For a universal constant $k > 0$, there is an explicit Hecke operator $\mathscr{T}$ with k generators on $\mathcal{L}^2(SO(n))$ with a spectral gap $1 - \lambda$ bounded away from 0.*

We now show how to obtain orthogonal t-designs using the above corollary. The idea itself is standard (see for example [45]) and we again use the intuition for finite graphs to motivate it: imagine running a random walk on the Cayley graph for a few ($\sim \log(1/\epsilon)$) steps. In the finite dimensional world, we expect that the resulting distribution on the vertices of the graph to be "close" to ($\sim \epsilon$) uniform.

In our setting, recall that our aim is to construct an object that fools the Haar ("uniform") distribution on $\mathsf{SO}(n)$. If we start a "random walk" on a Cayley graph with generators $g_1, g_2, \ldots, g_k$ on $\mathsf{SO}(n)$ from some fixed point, we expect the resulting distribution to be close to "uniform" on $\mathsf{SO}(n)$ after a few steps. This argument can be formalized to yield ϵ-approximate orthogonal 1-designs, i.e. those that fool all linear functions in the entries of the matrices drawn according to the Haar distribution on $\mathsf{SO}(n)$. To fool higher degree polynomials, we first take the tensor powers of the generators of the "Cayley graph" on $\mathsf{SO}(n)$. The entries of $g_i^{\otimes t}$, the t-wise tensor (Kronecker) product of g_i with itself, are all monomials of degree at most t in the entries of g_i. Thus if we start with the Cayley graph with the generators given by the t^{th} tensor powers of g_i, and argue similarly as above, we should hope to get approximate orthogonal t-designs.

We now show how to formalize this argument starting from the Hecke operator $\mathscr{T}$ given by Bourgain and Gamburd [9]. Suppose $g_1, g_2, \ldots, g_k$ are generators of the Hecke operator $\mathscr{T}$ as in the above corollary. Then, the constant function $\mathbb{1} : \mathsf{SO}(n) \to \mathbb{R}$, $\mathbb{1}(g) = 1$ every $g \in \mathsf{SO}(n)$ is an eigenfunction of T with eigenvalue 1. In particular, the operator $\mathscr{T}'$ defined by

$$\mathscr{T}'\rho \stackrel{\text{def}}{=} \mathscr{T}\rho - \mathbb{E}_{g \sim \mathbb{H}}[\rho(g)],$$

has a spectrum and all its eigenvalues are at most λ.

Define

$$S \stackrel{\text{def}}{=} 1/k \cdot \sum_{i=1}^{k} g_i^{\otimes t} - \mathbb{E}_{g\sim\mathbb{H}}[g^{\otimes t}].$$

Then, $S \in \mathbb{R}^{n^t \times n^t}$. We show that the spectral norm of S, $\|S\| \leq \lambda$.

Lemma 6.1. *Let $\mathscr{T}$ be a Hecke operator on $\mathcal{L}^2(\mathsf{SO}(n))$ with generators $g_1, g_2, \ldots, g_k$ with second largest eigenvalue λ. For any positive integer $t \in \mathbb{N}$, let $S = \left(\frac{1}{2k} \cdot \sum_{i=1}^{k} g_i^{\otimes t} - \mathbb{E}_{g\sim\mathbb{H}}[g^{\otimes t}]\right) \in \mathbb{R}^{n^t \times n^t}$. Then,*

$$\|S\| \leq \lambda.$$

Proof. Let $v \in \mathbb{R}^{n^t}$, $||v|| = 1$, be an eigenvector of S with eigenvalue θ. We will show that there is an eigenfunction $\gamma \in \mathcal{L}^2(\mathsf{SO}(n))$, $\gamma : \mathsf{SO}(n) \to \mathbb{R}$ of $\mathscr{T}'$ with the corresponding eigenvalue θ. Since $\|\mathscr{T}'\| \leq \lambda$, we will have the result of the lemma.

Define $\gamma(g) = \langle g^{\otimes t} \cdot v, e_1 \rangle$ where $e_1 \in \mathbb{R}^{n^t}$ with the first coordinate 1 and 0 otherwise. It is easy to observe that $\gamma \in \mathcal{L}^2(\mathsf{SO}(n))$. We now verify that γ is an eigenfunction of T'.

For any $h \in \mathsf{SO}(n)$,

$$\begin{aligned}
\mathscr{T}' \cdot \gamma(h) &= \frac{1}{2k} \cdot \sum_{i=1}^{k} (\mathscr{T}_{g_i} \cdot \gamma)(h) - \mathbb{E}_{g\sim\mathbb{H}}[\gamma(g)] \\
&= \frac{1}{2k} \cdot \sum_{i=1}^{k} \gamma(h \cdot g_i) - \mathbb{E}_{g\sim\mathbb{H}}[\gamma(h \cdot g)] \\
&= \frac{1}{2k} \cdot \sum_{i=1}^{k} \langle h^{\otimes t} \cdot g_i^{\otimes t} \cdot v, e_1 \rangle - \langle h^{\otimes t} \cdot \mathbb{E}_{g\sim\mathbb{H}}[g^{\otimes t}] \cdot v, e_1 \rangle \\
&= \langle h^{\otimes t} \cdot \frac{1}{2k} \cdot \sum_{i=1}^{k} g_i^{\otimes t} \cdot v, e_1 \rangle - \langle h^{\otimes t} \cdot \mathbb{E}_{g\sim\mathbb{H}}[g^{\otimes t}] \cdot v, e_1 \rangle \\
&= \langle h^{\otimes t} \cdot S \cdot v, e_1 \rangle = \theta \cdot \langle h^{\otimes t} \cdot v, e_1 \rangle = \theta \cdot \gamma(h),
\end{aligned}$$

□

where we use $(h \cdot g)^{\otimes t} = h^{\otimes t} \cdot g^{\otimes t}$ (which can be proven using induction and the mixed product property of the Kronecker product).

We can now use the result above to derive the main theorem of this section.

Proof of Lemma 5.1. Let $\mathscr{T}$ be a Hecke operator on $\mathcal{L}^2(\mathsf{SO}(n))$ with the second largest eigenvalue at most λ. Then, as above, $\mathscr{T}' : \mathcal{L}^2(\mathsf{SO}(n)) \to \mathcal{L}^2(\mathsf{SO}(n))$ defined by $\mathscr{T}' f(g) = \mathscr{T} f(g) - \mathbb{E}_{g\sim\mathbb{H}}[f]$ satisfies $\|\mathscr{T}'\| \leq \lambda$. Further, for any positive integer q, $\|\mathscr{T}'^q\| \leq \lambda^q$. We will choose q appropriately later.

Let $g_1, g_2, \ldots, g_k$ be the generators $\mathscr{T}^q$ and let $\mathcal{D}$ be a uniform draw from $\{g_1, g_2, \ldots, g_k\}$. We will show that $\mathcal{D}$ is an ϵ-approximate orthogonal t-design.

From Lemma 6.1 applied to $\mathscr{T}^q$, we know that $S = 1/k \cdot \sum_{i=1}^{k} g_i^{\otimes t} - \mathbb{E}_{g\sim\mathbb{H}}[g^{\otimes t}]$ satisfies $\|S\| \leq \lambda^q$.

We must show that for every polynomial $p : \mathsf{SO}(n) \to \mathbb{R}$ of degree at most t with $\|p\|_1 = 1$,

$$|\mathbb{E}_{g\sim\mathbb{H}}[p(g)] - \mathbb{E}_{g\sim\mathcal{D}}[p(g)]| \leq \frac{\epsilon}{n^t},$$

for some $q = t \log(n) + \log(1/\epsilon)$. Observe that it is enough to show this statement for monomials.

Let $M : \mathsf{SO}(n) \to \mathbb{R}$ be any monomial of degree $d \leq t$ (i.e. $M(g)$ is a product of at most t entries of the matrix g for any $g \in \mathsf{SO}(n)$). Then, $M(g)$ is an entry of the matrix $g^{\otimes d}$. Thus,

$$|\mathbb{E}_{g\sim\mathcal{D}}[M(g)] - \mathbb{E}_{g\sim\mathbb{H}}[M(g)]| \leq \left\| \frac{1}{k} \sum_{i=1}^{k} {g_i}^{\otimes d} - \mathbb{E}_{g\sim\mathbb{H}}[g^{\otimes d}] \right\| \leq \lambda^q.$$

Thus, choosing $q = \Theta(t \log(n) + \log(1/\epsilon))$ is enough. Thus, $\mathcal{D}$ is an ϵ-approximate orthogonal t-design.

□

Acknowledgment

We thank the anonymous reviewers for their suggestions on better presentation of the paper and pointing out the typos in a previous version.

References

[1] Matthias Krause 0001. Geometric arguments yield better bounds for threshold circuits and distributed computing. *Theor. Comput. Sci.*, 156:99–117, 1996.

[2] Matthias Krause 0001 and Stephan Waack. Variation ranks of communication matrices and lower bounds for depth two circuits having symmetric gates with unbounded fan-in. In *FOCS*, pages 777–782. IEEE Computer Society, 1991.

[3] N. I. Akhiezer. *The classical moment problem and some related questions in analysis*. Translated by N. Kemmer. Hafner Publishing Co., New York, 1965.

[4] A. Ambainis, J. Bouda, and A. Winter. Nonmalleable encryption of quantum information. *Journal of Mathematical Physics*, 50(4):042106, April 2009.

[5] Andris Ambainis and Joseph Emerson. Quantum t-designs: t-wise independence in the quantum world. In *22nd Annual IEEE Conference on Computational Complexity (CCC 2007), 13-16 June 2007, San Diego, California, USA*, pages 129–140, 2007.

[6] Animashree Anandkumar, Daniel Hsu, and Sham M. Kakade. A method of moments for mixture models and hidden markov models. In *COLT*, pages 33.1–33.34, 2012.

[7] Alexander Barvinok. Measure concentration. `http://www.math.lsa.umich.edu/~barvinok/total710.pdf`, 2005.

[8] Ido Ben-Eliezer, Shachar Lovett, and Ariel Yadin. Polynomial threshold functions: Structure, approximation and pseudorandomness. *CoRR*, abs/0911.3473, 2009.

[9] Jean Bourgain and Alexander Gamburd. Spectral gaps in SU(d). *C. R. Math. Acad. Sci. Paris*, 348(11-12):609–611, 2010.

[10] F. G. S. L. Brandao, A. W. Harrow, and M. Horodecki. Local random quantum circuits are approximate polynomial-designs. *ArXiv e-prints*, August 2012.

[11] Daniel Bump. *Lie groups*, volume 225 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2013.

[12] L. Elisa Celis, Omer Reingold, Gil Segev, and Udi Wieder. Balls and bins: Smaller hash families and faster evaluation. *SIAM J. Comput.*, 42(3):1030–1050, 2013.

[13] Bernard Chazelle. *The discrepancy method: randomness and complexity*. Cambridge University Press, New York, NY, USA, 2000.

[14] C. Daskalakis and C. Papadimitriou. Sparse Covers for Sums of Indicators. *ArXiv e-prints*, June 2013.

[15] Michael L. Dertouzos. Threshold logic: A synthesis approach. MIT Press, 1965.

[16] Ilias Diakonikolas, Parikshit Gopalan, Ragesh Jaiswal, Rocco A. Servedio, and Emanuele Viola. Bounded independence fools halfspaces. *SIAM J. Comput.*, 39(8):3441–3462, 2010.

[17] Ilias Diakonikolas, Daniel M. Kane, and Jelani Nelson. Bounded independence fools degree-2 threshold functions. In *FOCS*, pages 11–20, 2010.

[18] P. Dubey and L. S. Shapley. Mathematical properties of the banzhaf power index. Mathematics of Operations Research, 4:99-131, 1979.

[19] Jrgen Forster, Matthias Krause 0001, Satyanarayana V. Lokam, Rustam Mubarakzjanov, Niels Schmitt, and Hans-Ulrich Simon. Relations between communication complexity, linear arrangements, and computational complexity. In Ramesh Hariharan, Madhavan Mukund, and V. Vinay, editors, *FSTTCS*, volume 2245 of *Lecture Notes in Computer Science*, pages 171–182. Springer, 2001.

[20] Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.*, 55(1):119–139, 1997.

[21] Ju. R. Gabovič. The stability of the characterization of the multivariate normal distribution in the Skitovič-Darmois theorem. *Zap. Naučn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI)*, 61:5–16, 135, 1976. Continuity and stability in problems of probability theory and mathematical statistics.

[22] Mikael Goldmann, Johan Hstad, and Alexander A. Razborov. Majority gates vs. general weighted threshold gates. *Computational Complexity*, 2:277–300, 1992.

[23] Mikael Goldmann and Marek Karpinski. Simulating threshold circuits by majority circuits. *SIAM J. Comput.*, 27(1):230–246, 1998.

[24] Parikshit Gopalan, Daniel Kane, and Raghu Meka. Pseudorandomness for tail bounds and majorities, 2014. Manuscript.

[25] D. Gross and J. Eisert. Quantum Margulis expanders. *ArXiv e-prints*, October 2007.

[26] Andrs Hajnal, Wolfgang Maass, Pavel Pudlk, Mario Szegedy, and Gyrgy Turn. Threshold circuits of bounded depth. *J. Comput. Syst. Sci.*, 46(2):129–154, 1993.

[27] Aram W. Harrow. Quantum expanders from any classical cayley graph expander. *Quantum Info. Comput.*, 8(8):715–721, September 2008.

[28] Aram Wettroth Harrow and Richard A. Low. Efficient quantum tensor product expanders and k-designs. In *APPROX-RANDOM*, pages 548–561, 2009.

[29] Matthew B. Hastings and Aram Wettroth Harrow. Classical and quantum tensor product expanders. *Quantum Information & Computation*, 9(3):336–360, 2009.

[30] Szetsen Hu. Threshold logic. University of California, Berkeley, 1965.

[31] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. Pseudorandomness for network algorithms. In *STOC*, pages 356–364, 1994.

[32] J. R. Isbell. A counterexample in weighted majority games. 1969.

[33] A. V. Kakosyan, L. B. Klebanov, and S. T. Rachev. *Kolichestvennye kriterii skhodimosti veroyatnostnykh mer*. "Aĭastan", Erevan, 1988.

[34] Daniel Kane. A pseudorandom generator for polynomial threshold functions of gaussian with subpolynomial seed length. In *Computational Complexity (CCC), 2014 IEEE 29th Conference on*, pages 217–228, June 2014.

[35] Daniel M. Kane. k-independent gaussians fool polynomial threshold functions. In *IEEE Conference on Computational Complexity*, pages 252–261, 2011.

[36] Daniel M. Kane. A small prg for polynomial threshold functions of gaussians. In *FOCS*, pages 257–266, 2011.

[37] Daniel M. Kane. A structure theorem for poorly anticoncentrated gaussian chaoses and applications to the study of polynomial threshold functions. In *FOCS*, pages 91–100, 2012.

[38] Daniel M. Kane, Adam Klivans, and Raghu Meka. Learning halfspaces under log-concave densities: Polynomial approximations and moment matching. In *COLT*, pages 522–545, 2013.

[39] Zohar Shay Karnin, Yuval Rabani, and Amir Shpilka. Explicit dimension reduction and its applications. *SIAM J. Comput.*, 41(1):219–249, 2012.

[40] L. B. Klebanov and S. T. Mkrtčjan. Estimate of the closeness of distributions in terms of identical moments. In *Problems of stability of stochastic models (Proc. Fourth All-Union Sem., Palanga, 1979) (Russian)*, pages 64–72, 124. "Vsesojuz. Nauč.-Issled. Inst. Sistem. Issled.", Moscow, 1980.

[41] Henry J. Landau, editor. *Moments in mathematics*, volume 37 of *Proceedings of Symposia in Applied Mathematics*. American Mathematical Society, Providence, RI, 1987. Papers from the American Mathematical Society annual meeting held in San Antonio, Tex., January 20–22, 1987, AMS Short Course Lecture Notes.

[42] Jean B. Lasserre. Moments and positive polynomials for optimization. In *Modern optimization modelling techniques*, Adv. Courses Math. CRM Barcelona, pages 1–130. Birkhäuser/Springer Basel AG, Basel, 2012.

[43] M. Ledoux. *The Concentration of Measure Phenomenon*. Mathematical Surveys and Monographs. American Mathematical Society, 2005.

[44] Philip M. Lewis and Clarence Leroy Coates. Threshold logic. Wiley, 1967.

[45] A. Lubotzky, R. Phillips, and P. Sarnak. Hecke operators and distributing points on S^2. II. *Comm. Pure Appl. Math.*, 40(4):401–420, 1987.

[46] J. Marcinkiewicz and A. Zygmund. Marcinkiewicz-zygmund inequality. `http://en.wikipedia.org/wiki/Marcinkiewicz%E2%80%93Zygmund_inequality`.

[47] Józef Marcinkiewicz. *Collected papers*. Edited by Antoni Zygmund. With the collaboration of Stanislaw Lojasiewicz, Julian Musielak, Kazimierz Urbanik and Antoni Wiweger. Instytut Matematyczny Polskiej Akademii Nauk. Państwowe Wydawnictwo Naukowe, Warsaw, 1964.

[48] Daniel M. Kane Raghu Meka and Jelani Nelson. Almost optimal explicit Johnson-Lindenstrauss families. In *APPROX-RANDOM*, pages 628–639, 2011.

[49] Raghu Meka and David Zuckerman. Pseudorandom generators for polynomial threshold functions. *SIAM J. Comput.*, 42(3):1275–1301, 2013.

[50] Raghu Meka and David Zuckerman. Pseudorandom generators for polynomial threshold functions. *SIAM J. Comput.*, 42(3):1275–1301, 2013.

[51] S. Muroga. Threshold logic and its applications. Wiley-Interscience, New York, 1971.

[52] Noam Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.

[53] Frank Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.

[54] P. D. Seymour and Thomas Zaslavsky. Averaging sets: a generalization of mean values and spherical designs. *Adv. in Math.*, 52(3):213–240, 1984.

[55] Qinglai Sheng. Threshold logic. Ryerson Press, 1969.

[56] L. S.Penrose. The elementary statistics of majority voting. Journal of the Royal Statistical Society, Vol. 109, No. 1, 1946.

[57] Thomas Jan Stieltjes. Recherches sur les fractions continues. *Ann. Fac. Sci. Toulouse Math. (6)*, 4(4):A5–A47, 1995. Reprint of Ann. Fac. Sci. Toulouse **9** (1895), A5–A47.

[58] A. Taylor and W. Zwicker. A characterization of weighted voting. Proceedings of the AMS, 115(4):1089-1094, 1992.

[59] V. Vapnik. *Statistical Learning Theory*. Wiley-Interscience, New York, 1998.

[60] Roman Vershynin. Lectures on geometric functional analysis. 2011.

[61] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. In *Compressed sensing*, pages 210–268. Cambridge Univ. Press, Cambridge, 2012.

[62] Wikipedia. Faa di bruno's formula. `http://en.wikipedia.org/wiki/Fa%C3%A0_di_Bruno%27s_formula`.

A Deferred proofs

A.1 Proofs of facts from Section 4

Fact A.1 (Fact 4.2 restated)**.** *Let $g, h : \mathbb{R} \to \mathbb{R}$ be infinitely differentiable functions. Then,*

$$|(g \circ h)^{(k)}(x)| \leq (k!)^2 \cdot \left(\max_{\ell \leq k} |g^{(\ell)}(h(x))|\right) \cdot \left(\max_{\ell \leq k} |h^{(\ell)}(x)|\right)^k.$$

Proof. The result follows from using the following Faa di Bruno's formula ([62]) for derivatives of composition of functions (whenever all the derivatives in the expression exist):

$$\frac{d^m}{dt^m} g(h(t)) = \sum \frac{m!}{b_1! b_2! \ldots b_m!} g^{(\sum_{i=1}^m b_i)}(h(t)) \Pi_{i=1}^m \left(\frac{h^{(i)}}{i!}\right)^{b_i},$$

where the sum is over non-negative integers $b_1, b_2, \ldots, b_m$ such that $\sum_{i=1}^m i b_i = m$.

□

Fact A.2 (Fact 4.3 restated)**.** *For $h : (-1, \infty) \to \mathbb{R}$ be defined by $h(x) = 1/\sqrt{1+x}$. Then, $|h^{(k)}(x)| \leq k!$ for $-1/2 < x$.*

Proof.

$$h^{(k)}(x) = \frac{-1}{2} \cdot \frac{-3}{2} \cdots \frac{-(2k-1)}{2} \left(\frac{1}{1+x}\right)^{-(2k+1)/2}.$$

One can upper bound the expression on the RHS in absolute value by $k!$ for $x > -1/2$.

□

Fact A.3 (Fact 4.4 restated)**.** *Let $F : \mathbb{R} \to [0,1]$ be the CDF of a random variable Z. For a non-negative random variable V independent of Z, the CDF G of $V \cdot Z$ at any $t \in \mathbb{R}$ is given by:*

$$G(t) = \mathbb{E}_X[F(t/V)].$$

Proof. Let η be the PDF of V. We have:

$$\Pr[V \cdot Z \leq t] = \int_0^\infty Pr[Z \leq t/x] \cdot \eta(x) dx = \int_0^\infty F[t/x] \eta(x) dx = \mathbb{E}_V[F(t/V)].$$

□

Lemma A.1 (Lemma 3.2 restated)**.** *Let U, V_1, V_2 be independent random variables. Then,* $\mathsf{dcdf}(U \cdot V_1, U \cdot V_2) \leq \mathsf{dcdf}(V_1, V_2)$.

Proof. We can write the CDF of $U \cdot V_1$ as:

$$F_{UV_1}(t) = \Pr[UV_1 \leq t] = \int_0^\infty F_{V_1}(t/u) f_U(u) du + \int_{-\infty}^0 (1 - F_{V_1}(t/u) f_U(u) du.$$

Similarly, the CDF of $U \cdot V_2$ can be written as :

$$F_{UV_2}(t) = \int_0^\infty F_{V_2}(t/u) f_U(u) du + \int_{-\infty}^0 (1 - F_{V_2}(t/u) f_U(u) du.$$

Thus, $\mathsf{dcdf}(U \cdot V_1, U \cdot V_2) \leq \int_0^\infty \mathsf{dcdf}(V_1, V_2) f_U(u) du + \int_{-\infty}^0 \mathsf{dcdf}(V_1, V_2) f_U(u) du = \mathsf{dcdf}(V_1, V_2)$.

□

A.2 Proofs from Section 5

Recall that for the Gamma function Γ, $\Gamma(1/2) = \sqrt{\pi}$. The surface area of unit sphere $\mathbb{S}^{n-1}$ is $\frac{2\pi^{n/2}}{\Gamma(n/2)}$. We now obtain the PDF of the distribution of $< w, v >$ for $v \sim \mathbb{S}^{n-1}$. We now compute the PDF of $\langle w, v \rangle$ for $v \sim \mathbb{S}^{n-1}$.

Fact A.4. *Let $x = \langle w, v \rangle$ for v distributed uniformly over $\mathbb{S}^{n-1}$ and $w \in \mathbb{R}^n$, fixed, with $\|w\| = 1$. Then, x is supported on $[-1, 1]$ with the PDF:*

$$f(x) = \frac{1}{\sqrt{\pi}} \cdot \frac{\Gamma(n/2)}{\Gamma(\frac{n-1}{2})} \cdot (1 - x^2)^{\frac{n-3}{2}}.$$

Proof. Since the distribution of x is invariant under any rotation of both the vectors, we can assume that w has 1 in its first coordinate and 0 otherwise. Thus, $\langle w, v \rangle = v_1$. We can now calculate the CDF F of x by $F(x) = \Pr[v_1 \leq x]$.

Let $L_x \subseteq \mathbb{S}^{n-1}$ be defined by $L_x = \{z \in \mathbb{S}^{n-1} \mid z_1 \leq x\}$. Let $|L_x|$ denote the surface area of L_x. Then, $F(x) = |L_x|/|\mathbb{S}^{n-1}|$.

Let $t \in [-1, x]$. Then, $\{z \in L_x \mid z_1 = t\}$ defines a sphere of radius $\sqrt{1 - t^2}$ in $n - 2$ dimensions. Using the Jacobian of the area measure $1/\sqrt{1 - t^2}$, we can write $|L_x|$ as the integral:

$$|L_x| = \int_{-1}^{x} \frac{2\pi^{\frac{n-1}{2}}}{\Gamma(\frac{n-1}{2})} \cdot (1 - t^2)^{\frac{n-3}{2}} dt.$$

Thus,

$$F(x) = |L_x|/|\mathbb{S}^{n-1}| = 1/\sqrt{\pi} \cdot \int_{-1}^{x} \frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})} (1 - t^2)^{\frac{n-3}{2}} dt.$$

Now,

$$f(x) = F'(x) = \frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})} (1 - x^2)^{\frac{n-3}{2}}.$$

□

B Hecke operators with spectral gap on $\mathsf{SO}(n)$

In this section, we show that there exist Hecke operators on the group $\mathcal{L}^2(\mathsf{SO}(n))$ with a uniform spectral gap. This result follows almost immediately from the work of Bourgain and Gamburd [9], who show the existence of such operators on $\mathcal{L}^2(\mathsf{SU}(n))$. For completeness, we give a straightforward argument that uses only a few standard facts from the theory of Lie groups.

We state a few standard preliminary results (without proof) below. This material can be found in any standard textbook on Lie groups such as Bump [11].

B.1 Preliminaries

A *topological group* G is a group with an underlying topology on it such that the group operation $\cdot : G \times G \to G$ is a continuous map (with respect to the induced product topology on $G \times G$). *Lie groups* are topological groups where the group operation is *smooth*, that is, it has derivatives of all orders. Well known matrix groups, such as the General Linear group $GL_n(\mathbb{R})$: i.e. the

group of invertible $n \times n$ matrices on $\mathbb{R}$ and its subgroups, $SL_n(\mathbb{R})$: the subgroup of matrices with determinant 1 and $\mathsf{SO}(n)$: the subgroup of orthogonal matrices (referred to as the *rotation* group of the n-sphere) are all Lie groups. Similarly, the corresponding groups on the complex field $\mathbb{C}$: $GL_n(\mathbb{C})$, $SL_n(\mathbb{C})$ and $\mathsf{SU}(n)$ are also Lie groups.

By viewing $\mathbb{C}$ as a two dimensional vector space over $\mathbb{R}$, we observe that $\mathsf{SO}(n)$ is a subgroup of $\mathsf{SU}(n)$. On the other hand, by observing that $\mathbb{C}$-linear maps on $\mathbb{C}^n \simeq \mathbb{R}^{2n}$ are strict subsets of $\mathbb{R}$-linear maps on $\mathbb{R}^{2n}$ (under the vector space transformation from $\mathbb{C}$ to $\mathbb{R}$), we observe that $\mathsf{SU}(n)$ is a subgroup of $SO(2n)$.

Under the standard Euclidean topology, $\mathsf{SU}(n)$ and $\mathsf{SO}(n)$ are compact subgroups of $GL_{2n}(\mathbb{R})$. They are also closed, as is evident from the fact that $\mathsf{SU}(n)$ and $\mathsf{SO}(n)$ both can be defined as subgroups of matrices with certain polynomial equality constraints on their entries and are thus inverse images of closed sets under a continuous map.

B.1.1 Haar measure and linear operators on $\mathcal{L}^2(G)$

For compact Lie groups G such as $\mathsf{SU}(n)$ and $\mathsf{SO}(n)$, there exists a probability measure, known as the *Haar* measure, $\mathbb{H}$ (we also use dg to denote infinitesimal on G with respect to $\mathbb{H}$), on G that is invariant under (right or left) multiplication by any group element (this is written as the property of being (right or left) G-invariant). When both the right and left G-invariant probability measures coincide, G is said to be *unimodular*. It is a well known fact that both $\mathsf{SU}(n)$ and $\mathsf{SO}(n)$ are unimodular groups.

We can use the Haar measure to define integrals and (Hermitian) inner products of any two functions $p, q : G \to \mathbb{C}$:

$$\langle p, q \rangle = \int_{\mathbb{H}} p \cdot \bar{q} dg = \mathbb{E}_{\mathbb{H}}[p \cdot \bar{q}],$$

where $\bar{q}$ denotes the complex conjugate of q. Similarly, we can define the ℓ_2 norm of a function $f \in \mathcal{L}^2(G)$ by setting $\|\rho\| = \sqrt{\int_{\mathbb{H}} f \cdot \bar{f} dg} = \langle f, f \rangle$.

We can now define the linear space of square integrable functions on G: $\mathcal{L}^2(G) = \{f : G \to \mathbb{C} \mid \mathbb{E}_{g \sim \mathbb{H}}[f(g) \cdot \bar{f}(g)] < \infty\}$. This is a Hilbert space under the inner product defined above and forms, what is known as the *regular representation* of the group G under the (left or right) shift action. That is, for every $g \in G$, one can define a linear operator on $\mathcal{L}^2(G)$, T_g such that $(T_g f)(h) = f(h \cdot g)$ for every $h \in G$. Further, G as a group is *homomorphic* to the group of all such linear maps $\{T_g \mid g \in G\}$ under composition.

B.1.2 Push forward Haar measure on coset space of closed subgroups

Let G be a unimodular Lie group with a Haar measure $\mathbb{H}_G$ and let H be a closed unimodular subgroup of G with the Haar measure $\mathbb{H}_H$ defined on it. One can relate integrals of functions on either groups by constructing a G-invariant (probability) measure $\mathbb{H}_{G/H}$ on the coset space G/H. The existence of such a measure for compact unimodular Lie groups (a *push forward* measure) is a non-trivial but well known fact.

Let $f \in \mathcal{L}^2(G)$. Then, we have:

$$\int_{\mathbb{H}_G} f dg = \int_{\mathbb{H}_{G/H}} \int_{\mathbb{H}_H} f(h \cdot \dot{g}) dh d\dot{g},$$

where we use $\dot{g}$ to refer to the canonical element of the coset from G/H with $\dot{g}$ in it.

This, in particular lets us define a measure preserving embedding of $\mathcal{L}^2(H)$ into $\mathcal{L}^2(G)$. Let $\rho \in \mathcal{L}^2(H)$ with $\|\rho\|_2 = 1$. Choose canonical elements from G for every coset of H in G. If g belongs to the coset represented by the element $\dot{g} \in G$, then, $g = h \cdot \dot{g}$ for a unique element $h \in H$. We define $\tilde{\rho} : G \to \mathbb{C}$ by $\tilde{\rho}(g) = \rho(h)$ in this case. Observe that the construction is well defined as any two cosets of a subgroup inside a group are either disjoint or equal. We have thus created a map $f \to \tilde{f}$. It is easy to verify that the map is measure preserving using the relationship between the integrals on H and G given above:

$$\int_{\mathbb{H}_G} \tilde{\rho} dg = \int_{\mathbb{H}_{G/H}} \int_{\mathbb{H}_H} \tilde{\rho}(h \cdot \dot{g}) dh d\dot{g} = \int_{\mathbb{H}_{G/H}} \int_{\mathbb{H}_H} f(h) dh d\dot{g} = \int_{\mathbb{H}_H} \rho(h) dh.$$

In particular, $\|\tilde{\rho}\| = \|\rho\|$ and thus, the map defined above takes every $\rho \in \mathcal{L}^2(H)$ into a $\tilde{\rho} \in \mathcal{L}^2(G)$.

B.2 Existence of Hecke operators with spectral gap on $\mathsf{SO}(n)$

We are now ready to describe the existence of Hecke operators with spectral gap on $\mathsf{SO}(n)$. Bourgain and Gamburd [9] show the following:

Theorem 6 (Bourgain-Gamburd)**.** *For a universal constant $k > 0$, there is a Hecke operator with a spectral gap $\mathscr{T}$ on $\mathcal{L}^2(\mathsf{SU}(n))$ with k generators.*

Using the standard machinery developed in the preliminaries above the same result can be shown to hold for $\mathsf{SO}(n)$:

Corollary B.1 (Hecke Operators with Spectral Gap on $\mathsf{SO}(n)$)**.** *For a universal constant $k > 0$, there is a Hecke operator with a spectral gap $\mathscr{T}$ on $\mathcal{L}^2(\mathsf{SO}(n))$ with k generators.*

Proof. Recall that we see both $G = \mathsf{SU}(n)$ and $H = \mathsf{SO}(n)$ as *real* Lie groups by thinking of $\mathbb{C}$ as a two dimensional vector space over $\mathbb{R}$. Let $\mathscr{T}$ be the Hecke operator on $\mathcal{L}^2(G)$ with an inverse closed set of generators $g_1, g_2, \ldots, g_k \in SU(n)$ with a spectral gap given by the theorem above. That is, $\mathscr{T} = \frac{1}{k} \sum_{i=1}^{k} \mathscr{T}_{g_i}$.

The idea of the proof is the following: We want to define an operator $\mathscr{T}_h$ corresponding to T_g for any generator (or its inverse) g. We then want to argue that $\tilde{\mathscr{T}} : \mathcal{L}^2(H) \to \mathcal{L}^2(H)$ defined by $\tilde{\mathscr{T}} = \frac{1}{k} \sum_{i=1}^{k} \mathscr{T}_{h_i}$ where h_i correspond to g_i also has a spectral gap. To do this, we will use the embedding $\mathcal{E}$ of $\mathcal{L}^2(H)$ into $\mathcal{L}^2(G)$ defined above and if $\rho \to \tilde{\rho}$ is the embedding, then, we will show that $\tilde{\mathscr{T}}\tilde{\rho}$ corresponds to $\mathscr{T}'\rho$ in it. The spectral gap property for $\tilde{\mathscr{T}}$ will then follow from the spectral gap of $\mathscr{T}$.

We first define $\tilde{\mathscr{T}}$. Fix the canonical elements for every coset of H in G. Thus, $g = h \cdot \dot{g}$ whenever g belongs to the coset represented by $\dot{g}$. Let $h_g \in H$ be defined by $h_g = g \cdot \dot{g}^{-1}$. Set $\tilde{\mathscr{T}} = \frac{1}{k} \mathscr{T}_{h_{g_i}}$.

As before, let $\rho \in \mathcal{L}^2(H)$ such that $||\rho||_2 = 1$ and $\int_{\mathbb{H}_H} \rho(h) dh = 0$. Define $\tilde{\rho} \in \mathcal{L}^2(G)$ by $\tilde{\rho}(g) = \rho(g \cdot \dot{g}^{-1})$. Then from the discussion above, we know that $||\tilde{\rho}||_2 = 1$ and $\int_{\mathbb{H}_G} \tilde{\rho}(g) dg = 0$. Further, observe that $\mathscr{T}_g \tilde{\rho} = \mathcal{E}(\mathscr{T}_{h_g} \rho)$.

Define $Gen_G = \{g_1, g_2, \ldots, g_k\}$ as the (inverse closed) set of generators of $\mathscr{T}$ and let $Gen_H = \{h_g \mid g \in Gen_G\}$ be the corresponding set of generators for $\tilde{\mathscr{T}}$ as defined above.

We have:

$$\lambda^2 \geq \int_{\mathbb{H}_G} (\mathscr{T}\tilde{\rho})^2(g)dg = \frac{1}{2k} \sum_{g \in Gen_G} \int_{\mathbb{H}_{G/H}} \int_{\mathbb{H}_H} (\mathscr{T}_g \tilde{\rho})^2 (h \cdot \dot{g}) dh d\dot{g}$$
$$= \frac{1}{2k} \sum_{h_g \in Gen_H} \int_{\mathbb{H}_H} (\mathscr{T}_{h_g} \rho)^2 (h) dh = ||\tilde{\mathscr{T}}\rho||_2^2.$$

This completes the proof that $\tilde{\mathscr{T}}$ is a Hecke operator on $\mathcal{L}^2(H)$ with a spectral gap.

□